arXiv:1407.5770v1 [stat.CO] 22 Jul 2014

Perfect simulation using atomic regeneration with application to Sequential Monte Carlo

Anthony Lee*, Arnaud Doucet† and Krzysztof Łatuszyński*

July 23, 2014

*Department of Statistics, University of Warwick, UK

†Department of Statistics, University of Oxford, UK.

Abstract

Consider an irreducible, Harris recurrent Markov chain of transition kernel Π and invariant probability measure π. If Π satisfies a minorization condition, then the split chain allows the identification of regeneration times which may be exploited to obtain perfect samples from π. Unfortunately, many transition kernels associated with complex Markov chain Monte Carlo algorithms are analytically intractable, so establishing a minorization condition and simulating the split chain is challenging, if not impossible. For uniformly ergodic Markov chains with intractable transition kernels, we propose two efficient perfect simulation procedures of similar expected running time which are instances of the multigamma coupler and an imputation scheme. These algorithms overcome the intractability of the kernel by introducing an artificial atom and using a Bernoulli factory. We detail an application of these procedures when Π is the recently introduced iterated conditional Sequential Monte Carlo kernel. We additionally provide results on the general applicability of the methodology, and how Sequential Monte Carlo methods may be used to facilitate perfect simulation and/or unbiased estimation of expectations with respect to the stationary distribution of a non-uniformly ergodic Markov chain.

Keywords: Artificial atom; Bernoulli factory; Markov chain Monte Carlo; Perfect simulation; Regeneration; Sequential Monte Carlo.

1 Introduction

Given a target probability measure π on a general state space $(\mathsf{X}, \mathcal{B}(\mathsf{X}))$, the key idea of Markov chain Monte Carlo (MCMC) is to simulate a π-irreducible, Harris recurrent Markov chain $\mathbf{X} := (X_n)_{n\geq 1}$ with π-invariant transition kernel $\Pi : \mathsf{X}\times\mathcal{B}(\mathsf{X}) \to [0,1]$ to generate asymptotically samples from π. In practice, the Markov chain is only simulated for a finite number of iterations and so we do not obtain samples exactly distributed according to π. To address this problem, much effort has been put into the development of perfect simulation methods over the past twenty years. In particular, the Coupling From The Past (CFTP) procedure of Propp and Wilson (1996) has enabled the development of perfect simulation algorithms for the Ising model and various spatial point processes models (see, e.g., Møller and Waagepetersen, 2004). The applications of CFTP to general state spaces remain limited as, for implementation purposes, the MCMC kernel is typically required to satisfy strong stochastic monotonicity properties (see, e.g., Propp and Wilson, 1996; Foss and Tweedie, 1998).

We focus here on an alternative set of techniques based on regeneration. Henceforth, the first assumption we will make is that Π satisfies a one-step minorization condition; that is we have for all $x \in \mathsf{X}$

$$\Pi(x, \mathrm{d}y) \geq s(x)\nu(\mathrm{d}y), \tag{1}$$

where $s : \mathsf{X} \to [0, 1]$, ν is a given (regeneration) probability measure and $\pi(s) := \int_{\mathsf{X}} s(x)\pi(\mathrm{d}x) > 0$. In this context, the introduction of an associated "split chain" (Nummelin, 1978; Athreya and Ney, 1978) allows the identification of regeneration times and provides a mixture representation of π (Asmussen et al., 1992; Hobert and Robert, 2004).

This mixture representation was exploited by Hobert and Robert (2004) to obtain approximate samples from π, whereas Blanchet and Meng (2007) and Flegal and Herbei (2012) used it to obtain exact simulation algorithms but their procedures have an infinite expected running time. In the more restrictive scenario where one can take $s = \epsilon$ to be a constant function with $\epsilon > 0$, implying that Π is uniformly ergodic, the same mixture representation can be used to derive the multigamma coupler of Murdoch and Green (1998), an exact simulation algorithm with finite expected running time. Unfortunately, all of these techniques are fairly restrictive in the sense that they require being able to simulate the split chain, which can be difficult for sophisticated MCMC kernels. For example, the iterated conditional Sequential Monte Carlo (i-cSMC) kernel is an MCMC kernel introduced in Andrieu et al. (2010) to sample from high-dimensional target distributions. It has been established recently by Chopin and Singh (2013), Andrieu et al. (2013) and Lindsten et al. (2014) that this kernel satisfies the minorization condition (1) with $s = \epsilon > 0$ where ϵ can be known explicitly, yet one is unable to simulate the split chain associated to the i-cSMC kernel as this kernel does not admit a tractable expression.

The main contribution of this paper is to develop two general-purpose procedures for perfect simulation when the Markov transition kernel Π admits a singleton atom $\alpha = \{a\}$, for some $a \in \mathsf{X}$, for which $\inf_{x \in \mathsf{X}} \Pi(x, \alpha) \geq \beta > 0$ with β known. This, together with the ability to simulate the Markov chain $\mathbf{X}$, is the only requirement for perfect simulation to be implemented. While this assumes that $\mathbf{X}$ is uniformly ergodic, and knowledge of β is non-trivial in general, it is a significant relaxation of the conditions needed by other perfect simulation algorithms on general state spaces. The key mechanism that allows the identification of perfect samples is the use of a Bernoulli factory (Keane and O'Brien, 1994). We show that for specific implementations, the expected number of Markov chain transitions required to obtain a perfect sample is in $\mathcal{O}(\beta^{-1})$. Since many Markov transition kernels used in statistical applications do not admit a singleton atom, we overview relevant strategies in Brockwell and Kadane (2005) to modify Π to introduce an artificial singleton atom and obtain an associated invariant distribution which is a mixture of π and a point mass at this artificial atom. This modification can be carried out in a way which ensures that uniform ergodicity of the original Markov kernel is inherited by the modified Markov kernel.

While our methodology is generally applicable, we primarily focus in this paper on perfect simulation from the path distribution of a discrete-time Feynman–Kac model using an i-cSMC kernel. While it can be established that standard SMC methods provide samples whose distribution can be made arbitrary close to the path distribution of interest by increasing the number of particles, no perfect simulation method has hitherto been devised with guarantees of expected time polynomial in the time horizon, n. Under regularity assumptions on the Feynman–Kac model, we show that our methodology requires expected $\mathcal{O}(n^2)$ time to generate perfect samples, and may be implementable in expected polynomial time under less restrictive assumptions. The introduction of the artificial atom in this case follows from a simple and generally applicable extension of the Feynman–Kac model of interest. This allows us to sample from the joint posterior distribution of latent variables in a hidden Markov model (HMM), as long as the chosen parameter β satisfies $\inf_{x \in \mathsf{X}} \Pi(x, \alpha) \geq \beta$. In practice, we have found that with appropriate algorithm settings one can take β to be fairly large with no indication that this assumption fails to hold. Additionally a diagnostic can be performed to check this assumption during the course of the algorithm.

The rest of the paper is organized as follows. In Section 2 we review the split chain construction, its simulation and how this chain may be exploited to obtain perfect samples from π. In Section 3 we consider the case where Π admits a singleton atom, and show how perfect simulation methods

may be viably implemented through the use of an appropriate Bernoulli factory. In Section 4 we overview the general principles behind modification of Π to create a Markov kernel that admits an artificial singleton atom. Section 5 constitutes the major application, in which we combine results from Section 3 with a novel variant of an approach in Section 4 to develop perfect simulation methodology for sampling from a Feynman–Kac law on the path space. In Section 6 we discuss possible approaches for dealing with the case where the constant β is not known explicitly and show how the method we have introduced may be useful as a component of a more general perfect simulation algorithm or unbiased estimation scheme. We also reiterate the potential utility of parallel implementation of regenerative Markov chains. Section 7 demonstrates our methodology on a number of applications.

2 Regeneration and perfect simulation

The following notation will be used throughout the paper. For a metric space X, we denote by $\mathcal{B}(\mathsf{X})$ the Borel σ-algebra on X. For $\mu : \mathcal{B}(\mathsf{X}) \to \mathbb{R}_+$ a measure, $P : \mathsf{X} \times \mathcal{B}(\mathsf{X}) \to [0,1]$ a Markov kernel and f a real-valued measurable function w.r.t. $\mathcal{B}(\mathsf{X})$, we write $\mu(f) := \int_{\mathsf{X}} f(x)\mu(\mathrm{d}x)$ and define μP to be the measure satisfying $\mu P(\mathrm{d}y) := \int_{\mathsf{X}} P(x, \mathrm{d}y)\mu(\mathrm{d}x)$. We set $P^1 := P$ and $P^n(x, \mathrm{d}y) := \int_{\mathsf{X}} P^{n-1}(x, \mathrm{d}z)P(z, \mathrm{d}y)$ for $n \geq 2$. When a measure μ admits a density w.r.t. some dominating measure, we will denote both the density and the measure by μ, so e.g., $\mu(x)$ is the density at $x \in \mathsf{X}$ and $\mu(A)$ the measure of the set $A \in \mathcal{B}(\mathsf{X})$. When μ is additionally a probability measure, we will also refer to it as a distribution. If $f : \mathsf{X} \to \mathbb{R}$ is a function and $c \in \mathbb{R}$ a constant, we will write $f = c$ to mean f is the constant function with $x \mapsto c$ for all $x \in \mathsf{X}$. Finally, δ_x denotes the Dirac measure centred at x.

2.1 Atoms, split chain and regeneration

All forthcoming developments rely on the notion of an atom (see, e.g., Nummelin 1984, Definition 4.3 or Meyn and Tweedie 2009, Chapter 5), which we now introduce. Consider an irreducible, Harris recurrent Markov chain $\mathbf{Y} := (Y_n)_{n\geq 1}$ of transition kernel P of invariant distribution χ on a measurable space $(E, \mathcal{E})$. A set $\alpha \in \mathcal{E}$ is a *proper atom* for P if there exists a probability measure $\mu : \mathcal{E} \to [0,1]$ such that

$$P(y, A) = \mu(A), \quad y \in \alpha, A \in \mathcal{E},$$

that is each time the chain enters the atom, its next state is sampled according to the *regeneration measure* μ. Since $\mathbf{Y}$ is Harris recurrent and χ-irreducible, such an atom α is *accessible* if $\chi(\alpha) > 0$ and $\mathbf{Y}$ returns infinitely often to α with probability 1.

In general state spaces, it is rare for a Markov kernel to admit a proper, accessible atom. The major contribution of Nummelin (1978) and Athreya and Ney (1978) was to show that, even if the Markov chain $\mathbf{X}$ with transition kernel Π defined in Section 1 does not admit a proper atom, one can exploit the minorization (1) to construct a bivariate Markov chain $\tilde{\mathbf{X}}_{\nu,s}$ evolving on the extended space $\mathsf{X} \times \{0,1\}$ which admits a proper, accessible atom. We provide here a brief summary of the construction of Nummelin (1978).

Using (1), we can write $\Pi(x, \mathrm{d}y)$ as a mixture

$$\Pi(x, \mathrm{d}y) = s(x)\nu(\mathrm{d}y) + [1 - s(x)]R_{\nu,s}(x, \mathrm{d}y), \tag{2}$$

where $R_{\nu,s}$ is a *residual kernel*, defined for $x \in \mathsf{X}$ with $s(x) < 1$ by

$$R_{\nu,s}(x, \mathrm{d}y) := \frac{\Pi(x, \mathrm{d}y) - s(x)\nu(\mathrm{d}y)}{1 - s(x)}. \tag{3}$$

The subscript (ν, s) in $R_{\nu,s}$ emphasizes that there are many choices of ν and s in (1) for a given Π, each associated to a particular residual kernel. The split chain is then defined as follows.

Definition. The *split chain* is a bivariate Markov chain $\tilde{\mathbf{X}}_{\nu,s} = (\tilde{X}_n^{(\nu,s)})_{n\geq 1} = (Z_n^{(\nu,s)}, \rho_n^{(\nu,s)})_{n\geq 1}$ on the extended space $\mathsf{X} \times \{0,1\}$ with transition kernel

$$\tilde{\Pi}_{\nu,s}(x,\rho;\mathrm{d}y,\varrho) := \{\mathbb{I}(\rho=1)\nu(\mathrm{d}y) + \mathbb{I}(\rho=0)R_{\nu,s}(x,\mathrm{d}y)\}\, s(y)^{\varrho}\,[1-s(y)]^{1-\varrho}\,, \tag{4}$$

and invariant distribution $\tilde{\pi}_{\nu,s}(\mathrm{d}x,\rho) := \pi\,(\mathrm{d}x)\, s(x)^{\rho}\,[1-s(x)]^{1-\rho}$.

This construction, combined with the fact that $\mathbf{X}$ is Harris recurrent and π-irreducible, implies key properties of $\tilde{\mathbf{X}}_{\nu,s}$:

1. If the law of $Z_1^{(\nu,s)}$ is equal to the law of X_1 then the laws of $(Z_n^{(\nu,s)})_{n\geq 1}$ and $(X_n)_{n\geq 1}$ are identical (Nummelin, 1978, Theorem 1).
2. $\tilde{\mathbf{X}}_{\nu,s}$ is Harris recurrent with $\mathsf{X}\times\{1\}$ a proper, accessible atom for $\tilde{\Pi}_{\nu,s}$ with associated *regeneration measure* $\tilde{\nu}_{\nu,s}(\mathrm{d}x,\rho) := \nu(\mathrm{d}x)s(x)^{\rho}\,[1-s(x)]^{1-\rho}$ (Nummelin, 1978, Theorem 2).

Crucially, one can see that $\tilde{\Pi}_{\nu,s}$ admits $\mathsf{X}\times\{1\}$ as a proper, accessible atom even when the original Markov transition kernel Π admits no proper, accessible atom. To simplify notation and emphasize that only the second coordinate of $\tilde{\mathbf{X}}_{\nu,s} = (Z_n^{(\nu,s)}, \rho_n^{(\nu,s)})_{n\geq 1}$ is dependent on ν and s in (1), we dispense with $(Z_n^{(\nu,s)})_{n\geq 1}$ and hereafter define $\tilde{\mathbf{X}}_{\nu,s} := (\tilde{X}_n^{(\nu,s)})_{n\geq 1} := (X_n, \rho_n^{(\nu,s)})_{n\geq 1}$, the Markov chain with transition kernel $\tilde{\Pi}_{\nu,s}$. In this split chain, the variables $(\rho_n^{(\nu,s)})_{n\geq 1}$ are indicators of regeneration since $\tilde{X}_n \mid \{\rho_{n-1}^{(\nu,s)} = 1\} \sim \tilde{\nu}_{\nu,s}$.

When $\rho_{n-1}^{(\nu,s)} = 1$, it is customary to call $\tilde{X}_{n-1}^{(\nu,s)}$ the sample just prior to regeneration, since it is $\tilde{X}_n^{(\nu,s)}$ that is distributed according to $\tilde{\nu}_{\nu,s}$. Due to the close relationship between $\mathbf{X}$ and $\tilde{\mathbf{X}}_{\nu,s}$, ν is commonly also referred to as a regeneration measure for $\mathbf{X}$, the construction of an associated split chain $\tilde{\mathbf{X}}_{\nu,s}$ being implicit.

We define the sequence of regeneration times $(\tau_k^{(\nu,s)})_{k\geq 1}$ of the Markov chain $\tilde{\mathbf{X}}_{\nu,s}$ via

$$\tau_k^{(\nu,s)} := \min\{n > \tau_{k-1}^{(\nu,s)} : \rho_n^{(\nu,s)} = 1\},$$

with $\tau_0^{(\nu,s)} := 0$. In the sequel, $\tau_{\nu,s} := \tau_1^{(\nu,s)}$ will be used to denote the first regeneration time where no ambiguity can result. The introduction of this split chain was a major innovation in the analysis of Markov chains on general state spaces, because it allows the Markov chain $\tilde{\mathbf{X}}_{\nu,s}$ to be partitioned into i.i.d. tours, where tour i, $i \in \mathbb{N}$, is defined as $(\tilde{X}_n^{(\nu,s)} : \tau_{i-1}^{(\nu,s)} < n \leq \tau_i^{(\nu,s)})$.

2.2 Simulating the split chain

Although the split chain $\tilde{\mathbf{X}}_{\nu,s}$ was originally introduced as a theoretical tool to analyze the Markov chain $\mathbf{X}$, statistical methodology has since been developed which requires simulating $\tilde{\mathbf{X}}_{\nu,s}$ (see, e.g., Mykland et al., 1995; Hobert et al., 2002). In practice, however, sampling from the split chain kernel $\tilde{\Pi}$ is not always feasible even when sampling from Π is.

A standard approach due to Mykland et al. (1995) consists of simulating $\mathbf{X}$ and then imputing the values of $(\rho_n^{(\nu,s)})_{n\geq 1}$ conditional upon $\mathbf{X}$. Letting $\mathbb{P}_{\nu,s}$ denote the law of $\tilde{\mathbf{X}}_{\nu,s}$ when $\tilde{X}_1 \sim \tilde{\nu}_{\nu,s}$, they observe that $\rho_{n-1}^{(\nu,s)}$ depends on $\tilde{\mathbf{X}}_{\nu,s}$ only through X_{n-1} and X_n and that

$$\mathbb{P}_{\nu,s}\left(\rho_{n-1}^{(\nu,s)} = 1 \mid X_{n-1} = x_{n-1}, X_n = x_n\right) = \frac{s(x_{n-1})\nu(\mathrm{d}x_n)}{\Pi(x_{n-1},\mathrm{d}x_n)}, \tag{5}$$

which can be computed in a variety of situations, but not in general. For example, in many cases Π can be expressed as

$$\Pi(x,\mathrm{d}y) = Q_x(\mathrm{d}y)a(x,y) + r(x)\delta_x(\mathrm{d}y),$$

where $a : \mathsf{X} \times \mathsf{Y} \to [0,1]$ and $r : \mathsf{X} \to [0,1]$, and $\{Q_x : x \in \mathsf{X}\}$ and ν admit densities with respect to a common dominating measure without point masses. Then $\mathbb{P}_{\nu,s}(\rho_{n-1}^{(\nu,s)} = 1 \mid X_{n-1} = X_n) = 0$ and for any $X_{n-1} \neq X_n$ the r.h.s. of (5) is given by (Mykland et al., 1995, Proposition 1)

$$\frac{s(x_{n-1})\nu(x_n)}{Q_{x_{n-1}}(x_n)a(x_{n-1},x_n)}.$$

It is apparent, however, that imputing $\rho_{n-1}^{(\nu,s)}$ in practice may be impossible when this expression cannot be computed for distinct $x_{n-1}, x_n \in \mathsf{X}$. We will show in Section 3 how we can overcome this limitation when Π admits a singleton atom.

2.3 Mixture representation of the invariant measure and perfect simulation

It was established in Asmussen et al. (1992), Hobert and Robert (2004) and Hobert et al. (2006) that the split chain construction provides the following mixture representation of the target distribution π:

$$\pi(\mathrm{d}x) = \sum_{n=1}^{\infty} \frac{\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})} \eta_n^{(\nu,s)}(\mathrm{d}x), \tag{6}$$

where, defining $G_s(x) := 1 - s(x)$ and with the convention that $\prod_{\emptyset} = 1$,

$$\begin{aligned} \eta_n^{(\nu,s)}(A) &= \mathbb{P}_{\nu,s}\left(X_n \in A \mid \tau_{\nu,s} \geq n\right) \\ &= \frac{\int_{\mathsf{X}^n} \mathbb{I}\{x_n \in A\} \left[\prod_{p=2}^n G_s(x_{p-1})\right] \nu(\mathrm{d}x_1) \prod_{p=2}^n R_{\nu,s}(x_{p-1}, \mathrm{d}x_p)}{\int_{\mathsf{X}^n} \left[\prod_{p=2}^n G_s(x_{p-1})\right] \nu(\mathrm{d}x_1) \prod_{p=2}^n R_{\nu,s}(x_{p-1}, \mathrm{d}x_p)}. \end{aligned} \tag{7}$$

This mixture representation (6) implies that if one were able to sample a $\mathbb{N}$-valued random variable N of probability mass function (p.m.f.)

$$\Pr(N = n) = \frac{\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})} \tag{8}$$

and $\xi | (N = n) \sim \eta_n^{(\nu,s)}$ then unconditionally $\xi \sim \pi$. Algorithm 1 describes this procedure.

Algorithm 1 A generic regenerative perfect simulation algorithm

1. Sample N from the distribution with p.m.f. (8).
2. Output a sample from $\eta_N^{(\nu,s)}$.

It has been realized in Blanchet and Meng (2007) and Flegal and Herbei (2012) that a Bernoulli factory can be used in some circumstances to facilitate sampling N according to (8), although we note that the recent literature on the Bernoulli factory starting with Keane and O'Brien (1994) was itself motivated by the regenerative simulation work of Asmussen et al. (1992). For practical implementation, this typically requires accurate upper bounds on $\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)$ as a function of n, and Flegal and Herbei (2012) provide examples where this is possible. Sampling from $\eta_N^{(\nu,s)}$, and thereby obtaining a sample from π, is then achieved by using a simple rejection algorithm where the proposal is $\nu R_{\nu,s}^{N-1}$. Unfortunately, this rejection technique has an infinite expected running time when N is distributed according to (8) (Asmussen et al., 1992; Blanchet and Meng, 2007, Proposition 2).

As noted in Hobert and Robert (2004), the problem becomes much simpler when one considers the case where (1) holds with $s = \epsilon$. The representation (6) adapted to the minorization (1) with $s = \epsilon$ yields

$$\pi(\mathrm{d}x) = \sum_{n=1}^{\infty} \epsilon(1-\epsilon)^{n-1} \nu R_{\nu,\epsilon}^{n-1}(\mathrm{d}x). \tag{9}$$

Indeed, $\tau_{\nu,\epsilon}$ is a geometric random variable with success probability ϵ so

$$\frac{\mathbb{P}_{\nu,\epsilon}(\tau_{\nu,\epsilon} \geq n)}{\mathbb{E}_{\nu,\epsilon}(\tau_{\nu,\epsilon})} = \epsilon(1-\epsilon)^{n-1} = \mathbb{P}_{\nu,\epsilon}(\tau_{\nu,\epsilon} = n),$$

and $\eta_n^{(\nu,\epsilon)} = \nu R_{\nu,\epsilon}^{n-1}$ because $s(x) = \epsilon$ is independent of x. There are two perfect simulation procedures that arise from alternative interpretations of (9), which we now outline.

The first interpretation is that if $\tilde{X}_1 \sim \tilde{\nu}_{\nu,\epsilon}$ then $X_{\tau_{\nu,\epsilon}}$, the first coordinate of $\tilde{\mathbf{X}}_{\nu,\epsilon}$ at time $\tau_{\nu,\epsilon}$, is a perfect sample from π. In other words, the sample $X_{\tau_{\nu,\epsilon}}$ just prior to regeneration is a perfect sample from π, and Algorithm 2 is a corresponding perfect simulation procedure using (5) to determine when the first regeneration has occurred. The algorithm is practical whenever ϵ is known and one can compute the Radon–Nikodym derivative appearing in (10).

Algorithm 2 Regenerative perfect simulation algorithm when $s = \epsilon$ via imputation

1. Sample $X_1 \sim \nu$.
2. For $n = 2, 3, \ldots$
   (a) Sample $X_n \sim \Pi(X_{n-1}, \cdot)$.
   (b) With probability
   $$\epsilon \frac{\mathrm{d}\nu(\cdot)}{\mathrm{d}\Pi(X_{n-1}, \cdot)}(X_n), \tag{10}$$
   stop and output X_{n-1}.

A second interpretation is that if one samples $N \sim$ Geometric(ϵ) followed by $\xi \sim \nu R_{\nu,\epsilon}^{N-1}$, then unconditionally $\xi \sim \pi$, and Algorithm 3 is a procedure based upon this observation. The algorithm is practical as soon as ϵ is known and both ν and $R_{\nu,\epsilon}$ can be sampled from. This latter scheme corresponds to the multigamma coupler of Murdoch and Green (1998) but without any explicit appeal to CFTP as noticed by Hobert and Robert (2004).

Algorithm 3 Regenerative perfect simulation algorithm when $s = \epsilon$ via the multigamma coupler

1. Sample $N \sim$ Geometric(ϵ).
2. Output a sample from $\nu R_{\nu,\epsilon}^{N-1}$.

While cast as perfect simulation algorithms, both of these procedures essentially simulate a single tour of the split chain $\tilde{\mathbf{X}}_{\nu,\epsilon}$. In practice, it is often the case that one cannot sample from ν or $R_{\nu,\epsilon}$, or compute the Radon-Nikodym derivative (10) so that Algorithms 2 and 3 cannot be implemented. We now discuss a special case where such an implementation is feasible.

3 Regeneration and perfect simulation for Markov chains with a singleton atom

3.1 Markov chains with singleton atoms

Assume that the Markov chain $\mathbf{X}$ admits not only a proper, accessible atom α, but one which is additionally a *singleton*, i.e., $\alpha = \{a\}$ for some distinguished $a \in \mathsf{X}$ assumed known to the user. This special case is a central focus of this paper. Most Markov kernels do not admit such a singleton atom but we will show in Section 4 how such an atom can be introduced fairly generally,

following Brockwell and Kadane (2005). When this assumption is met, $\mathbf{X}$ will visit the point a infinitely often.

A natural split chain to consider in this context is obtained by taking $s = p$ and $\nu = \delta_a$ in (1) where

$$p(x) := \Pi(x, \{a\}).$$

To simplify notation, we write $\tilde{\mathbf{X}}_{a,p}$, $\rho_n^{(a,p)}$ and $\mathbb{P}_{a,p}$ for $\tilde{\mathbf{X}}_{\delta_a,p}$, $\rho_n^{(\delta_a,p)}$ and $\mathbb{P}_{\delta_a,p}$ respectively. The split chain $\tilde{\mathbf{X}}_{a,p}$ can be very easily simulated in this case using the imputation method discussed in Section 2.2 as (5) reduces to

$$\mathbb{P}_{a,p}\left(\rho_{n-1}^{(a,p)} = 1 \mid X_{n-1} = x_{n-1}, X_n = x_n\right) = \mathbb{I}\left(X_n = a\right).$$

However, performing perfect simulation using Algorithm 1 in this scenario remains challenging as all currently available implementations have an infinite expected running time (Asmussen et al., 1992; Blanchet and Meng, 2007; Flegal and Herbei, 2012). We show how to bypass this problem when $p(x)$ satisfies an additional assumption.

3.2 Perfect simulation using the split chain $\tilde{\mathbf{X}}_{a,\epsilon}$

We develop here practical perfect simulation schemes under the additional assumption that for some known constant $\beta > 0$

$$\underline{p} := \inf_{x \in \mathsf{X}} p(x) \geq \beta.$$

This assumption implies that $\mathbf{X}$ is uniformly ergodic and satisfies the one-step minorization condition $\Pi(x, \{a\}) \geq \beta$. The split chain $\tilde{\mathbf{X}}_{a,\epsilon}$ obtained by taking $s = \epsilon$ and $\nu = \delta_a$ in (1) for some $\epsilon \in (0, \beta)$ is at the heart of our methodology. Contrary to the natural split chain $\tilde{\mathbf{X}}_{a,p}$ introduced in Section 3.1 which does not lead to a currently implementable perfect simulation algorithm enjoying a finite expected running time, the introduction of the alternative split chain $\tilde{\mathbf{X}}_{a,\epsilon}$ is motivated by Section 2.3 where it is shown that perfect simulation from π is feasible if either Algorithm 2 or Algorithm 3 can be implemented.

We show here how one can implement both Algorithms 2 and 3 using an appropriate Bernoulli factory. Our algorithms exploit the fact that although the expression of $p(x)$ is usually unknown, one can simulate a Bernoulli random variable with success probability $p(x)$, for any required $x \in \mathsf{X}$, by sampling $Y \sim \Pi(x, \cdot)$ and outputting $\mathbb{I}\{Y = a\}$. We will generically refer to a p-coin as a coin whose flips are independent Bernoulli random variables with success probability p. Details of the solutions to the Bernoulli factory problems are presented in Section 3.3.

Algorithm 2 involves two operations: simulation from Π and simulation of a Bernoulli random variable with success probability (10). The former is possible by assumption and the latter reduces in this case to

$$\mathbb{P}_{a,\epsilon}\left(\rho_{n-1}^{(a,\epsilon)} = 1 \mid X_{n-1} = x_{n-1}, X_n = x_n\right) = \frac{\epsilon}{p(x_{n-1})}\mathbb{I}\{x_n = a\}, \tag{11}$$

which is possible as we have a solution to the following Bernoulli factory problem.

Problem 1. Given flips of a p-coin with $p \geq \beta > \epsilon > 0$, for known constants ϵ and β, simulate an $\frac{\epsilon}{p}$-coin flip.

The practical implementation of Algorithm 2 is summarized in Algorithm 4.

Algorithm 4 Practical implementation of Algorithm 2

1. Choose $\beta \in (0, \underline{p}]$ and $\epsilon \in (0, \beta)$. We suggest $\epsilon = \beta/2$.
2. Set $X_1 = a$.
3. For $n = 2, 3, \ldots$:
   (a) Sample $X_n \sim \Pi(X_{n-1}, \cdot)$.
   (b) If $X_n = a$, sample $\rho_{n-1}^{(a,\epsilon)} \sim \text{Bernoulli}(\epsilon/p(X_{n-1}))$. Otherwise set $\rho_{n-1}^{(a,\epsilon)} = 0$.
   (c) If $\rho_{n-1}^{(a,\epsilon)} = 1$, stop and output X_{n-1}.

The main computational task in Algorithm 3 is simulation from the residual kernel $R_{a,\epsilon}$, which can be decomposed into a mixture of $R_{a,p}$ and $\nu = \delta_a$

$$\begin{aligned} R_{a,\epsilon}(x, \mathrm{d}y) &= \frac{\Pi(x, \mathrm{d}y) - \epsilon\delta_a(\mathrm{d}y)}{1-\epsilon} \\ &= \frac{1-p(x)}{1-\epsilon} R_{a,p}(x, \mathrm{d}y) + \frac{p(x)-\epsilon}{1-\epsilon}\delta_a(\mathrm{d}y), \end{aligned} \tag{12}$$

where we can simulate from $R_{a,p}(x, \cdot)$ by a rejection method which simulates from the proposal $\Pi(x, \cdot)$ and outputs the first sample distinct from a. The difficulty in simulating from $R_{a,\epsilon}(x, \cdot)$ is in flipping a $\{1-p(x)\}/(1-\epsilon)$-coin to select the mixture component. However, this is possible as we have a solution to the following second Bernoulli factory problem.

Problem 2. Given flips of a p-coin with $p \geq \beta > \epsilon > 0$, for known constants ϵ and β, simulate a $(1-p)/(1-\epsilon)$-coin flip.

The practical implementation of Algorithm 3 is summarized in Algorithm 5.

Algorithm 5 Practical implementation of Algorithm 3

1. Choose $\beta \in (0, \underline{p}]$ and $\epsilon \in (0, \beta)$. We suggest $\epsilon = \beta/2$.
2. Sample $N \sim \text{Geometric}(\epsilon)$.
3. Set $X_1 = a$.
4. For $n = 2, \ldots, N$:
   (a) Sample $Y_n \sim \text{Bernoulli}\left(\{1-p(X_{n-1})\}/\{1-\epsilon\}\right)$.
   (b) If $Y_n = 1$, sample $X_n \sim R_{a,p}(X_{n-1}, \cdot)$ by rejection. Otherwise, set $X_n = a$.
5. Output X_N.

3.3 Bernoulli factory algorithms

A Bernoulli factory for a known function f is an algorithm for simulating a flip of a $f(p)$-coin when one can flip multiple times a p-coin, but p is unknown (Keane and O'Brien, 1994). More specifically, letting $\mathcal{P} \subseteq [0,1]$ and $f : \mathcal{P} \to [0,1]$, such a factory must output a flip of a $f(p)$-coin for any $p \in \mathcal{P}$ without knowledge of p but is allowed to flip a p-coin an almost surely finite number of times during its execution. One early example of such a factory was presented in Von Neumann (1951), where $f(p) = 1/2$. More recent interest arises from the perfect simulation of general regenerative processes (Asmussen et al., 1992). This inspired Keane and O'Brien (1994),

whose major contribution is a necessary and sufficient condition for the existence of a Bernoulli factory for f: that f is either constant, or continuous and satisfies, for some $n \geq 1$,

$$\min\{f(p), 1-f(p)\} \geq \min\{p, 1-p\}^n, \quad \forall p \in \mathcal{P}.$$

It can be verified from this result that Bernoulli factories exist for the functions given in Problems 1 and 2 when $p \geq \beta > \epsilon > 0$ for known constants β and ϵ.

While the existence of Bernoulli factories is shown in Keane and O'Brien (1994), the proof is not constructive. Bernoulli factory algorithms have been provided in a series of papers (Nacu and Peres, 2005; Łatuszyński et al., 2011; Thomas and Blanchet, 2011; Flegal and Herbei, 2012; Huber, 2014). In these, most attention is paid to the Bernoulli factory for f satisfying

$$cp \leq \gamma \implies f(p) = cp \tag{13}$$

for a given $c > 0$ and $\gamma \in (0, 1)$. The values that f takes when $cp > \gamma$ can be thought of as an "extension", particularly when it is known that the factory will only be invoked for a p-coin satisfying $cp \leq \gamma$. For example, in Nacu and Peres (2005) and Łatuszyński et al. (2011) the function $f(p) = \min\{\gamma, cp\}$ is treated, whose extension is thus the constant function γ. Flegal and Herbei (2012) propose an alternative extension so that f is twice differentiable, which provides some performance guarantees (Nacu and Peres, 2005, Proposition 10, see also Holtz et al., 2011). In Huber (2014), an algorithm is presented for which the extension is not explicitly shown, and which requires an expected number of flips of a p-coin which is bounded above by $9.5c/(1-\gamma)$ whenever $cp < \gamma$ (Huber, 2014, Theorem 1). In this paper all our algorithms are implemented using the efficient procedure in Huber (2014).

Problem 2 can be solved by a standard solution for f satisfying (13); in Appendix C we develop an alternative solution to Problem 2 via a generic solution to the sign problem, which may be of independent interest. We now show that Problem 1 can be solved using any solution to Problem 2.

Proposition 1. *One can flip an ϵ/p-coin where $p > \epsilon > 0$ by sampling $K \sim$ Geometric(ϵ) and flipping a $\{(1-p)/(1-\epsilon)\}^{K-1}$-coin.*

Algorithm 6 is an implementation of the procedure in Proposition 1, where a $\{(1-p)/(1-\epsilon)\}^{K-1}$-coin flip is viewed as the product of $K-1$ independent $(1-p)/(1-\epsilon)$-coin flips. The procedure can terminate as soon as any of these $(1-p)/(1-\epsilon)$-coin flips is 0. We have therefore formulated Algorithm 6 as a "race" between an ϵ-coin and a $(p-\epsilon)/(1-\epsilon)$-coin. One can indeed directly check that in this algorithm, the probability that $Y_n = 1$ conditional upon $Y_n + Z_n \geq 1$ is ϵ/p. We note that an alternative solution to Problem 1 is to simulate $K \sim$ Geometric(ϵ) and then simulate a $c_{K-1}\tilde{p}_{K-1}$-coin where $c_k = (1-\epsilon)^{-k}$ and $\tilde{p}_k = (1-p)^k$, but we do not pursue this further.

Algorithm 6 Simulate an ϵ/p-coin flip using a $(1-p)/(1-\epsilon)$-coin

For $n = 1, 2, \ldots$:

1. Simulate an ϵ-coin flip, Y_n. If $Y_n = 1$, stop and output 1.
2. Simulate a $(p-\epsilon)/(1-\epsilon)$-coin flip, Z_n. If $Z_n = 1$, stop and output 0.

3.4 Computational cost of algorithms 4 and 5

A striking relationship between Algorithms 4 and 5 is that they have almost exactly the same expected computational effort, when the ϵ/p-coin flips in Algorithm 4 are simulated using Algorithm 6.

Proposition 2. *The expected number of $(1-p)/(1-\epsilon)$-coins required to simulate an ϵ/p-coin using Algorithm 6 is $(1-\epsilon)/p$.*

Corollary 1. *The expected number of* $(1-p)/(1-\epsilon)$*-coins required to simulate* X_n *and* $\rho_{n-1}^{(a,\epsilon)}$ *in Algorithm 4 is* $1-\epsilon$*.*

Proposition 3. *Using either Algorithm 4 or 5, the expected number of* $(1-p)/(1-\epsilon)$*-coin flips required to simulate a single tour of the split chain* $\check{\mathbf{X}}_{a,\epsilon}$ *is* $\epsilon^{-1}-1$*, and the average number of samples from* Π *to additionally simulate the tour itself is* ϵ^{-1}*.*

Solving Problem 2 directly using the algorithm of Huber (2014) involves taking $c=(1-\epsilon)^{-1}$ and $\gamma=(1-\beta)/(1-\epsilon)$ in (13). We suggest that one takes $\beta \leq 0.5$ and $\epsilon=\beta/2$ because of the following result, which indicates that the expected number of p-coin flips required is then upper bounded by a small constant using the default algorithm settings in Huber (2014).

Proposition 4. *Let* $\beta \leq 0.5$*,* $\epsilon=\beta/2$ *and* $p \geq \beta$*. The expected number of* p*-coin flips to produce a* $(1-p)/(1-\epsilon)$*-coin flip using the algorithm of Huber (2014) with its default settings is bounded above by* 11*.*

As a result, when $\beta \leq 0.5$ and $\epsilon=\beta/2$, one is ensured that the expected cost of obtaining a perfect sample is bounded above by $12/\epsilon$. Empirically we have observed that this bound is fairly tight, but that the expected number of p-coin flips required to produce a $(1-p)/(1-\epsilon)$-coin flip may even be inferior to 7.

4 Introduction of an artificial singleton atom

4.1 State space extension

On general state spaces, the existence of accessible, singleton atoms is not guaranteed. In most statistical applications, for example, one has $\mathsf{X}=\mathbb{R}^d$ and π admits a density w.r.t. the Lebesgue measure, in which case no such atom for Π can exist since $\pi(\{a\})=0$ for any $a \in \mathsf{X}$. Brockwell and Kadane (2005) suggested in such cases to define a different transition kernel $\check{\Pi}$ on an extended state space for which such an *artificial* singleton atom exists. The atom is artificial in the sense that one defines an extended state space $\check{\mathsf{X}} := \mathsf{X} \cup \{a\}$ for some distinguished a, and seeks to design $\check{\Pi}$ such that it is Harris recurrent and irreducible with unique invariant probability measure $\check{\pi} : \mathcal{B}(\check{\mathsf{X}}) \to [0,1]$ satisfying, for some $k \in (0,1)$,

$$\check{\pi}(A) = k\pi(A) + (1-k)\mathbb{I}(a \in A), \quad A \in \mathcal{B}(\check{\mathsf{X}}). \tag{14}$$

When (14) holds, it follows that $\check{\pi}(A)=k\pi(A)$ for any $A \in \mathcal{B}(\mathsf{X})$. We denote by $\check{\mathbf{X}} := (\check{X}_n)_{n\geq 1}$ the Markov chain with transition kernel $\check{\Pi}$.

It is always possible to recover an irreducible, Harris recurrent Markov chain with invariant probability measure π from $\check{\mathbf{X}}$. If we define $\check{\mathbf{Y}} := (\check{X}_n)_{n\in\mathcal{J}}$ with $\mathcal{J} := \{n : \check{X}_n \in \mathsf{X}\}$ then $\check{\mathbf{Y}}$ is the Markov chain $\check{\mathbf{X}}$ "watched" in the set X and is irreducible and Harris recurrent with invariant probability measure π (Brockwell and Kadane, 2005, Theorem 1). One can also check that $\check{\mathbf{Y}}$ satisfies (1) with $s(x)=\check{\Pi}(x,\{a\})$ and $\nu(\mathrm{d}y)=\check{\Pi}(a,\mathrm{d}y)\mathbb{I}(y \in \mathsf{X})/\check{\Pi}(a,\mathsf{X})$ (Brockwell and Kadane, 2005, Theorem 2).

4.2 Practical design of the distribution $\check{\pi}$ and transition kernel $\check{\Pi}$

We give here practical ways to define the distribution $\check{\pi}$ and kernel $\check{\Pi}$ given π and Π. In many statistics applications π admits a density w.r.t. to a dominating measure λ on X and we can compute an unnormalized version $\gamma(x)$ of this density. In this case, we can choose a $b>0$ and define an unnormalized version $\check{\gamma}(x)$ of the density of $\check{\pi}$ w.r.t. the dominating measure $\lambda+\delta_a$ on $\check{\mathsf{X}}$ through

$$\check{\gamma}(x) := \mathbb{I}(x \in \mathsf{X})\gamma(x) + \mathbb{I}(x=a)b. \tag{15}$$

It follows from (15) that $\check{\pi}(\mathrm{d}x) = \check{\gamma}(x)\{\lambda(\mathrm{d}x) + \delta_a(\mathrm{d}x)\}/\check{\gamma}(\mathsf{X})$ satisfies (14) with $k = \{1 + b/\gamma(\mathsf{X})\}^{-1}$. In practice, it is desirable that $\check{\pi}(\{a\})$ be not too close to either 0 or 1 so that visits to the artificial atom are frequent and yet the Markov chain spends a substantial amount of time in the set X. This suggests that an estimate of $\gamma(\mathsf{X})$ is necessary to be able to choose an appropriate value of b.

Two strategies for designing $\check{\Pi}$ are proposed in Brockwell and Kadane (2005). The first one is to let $\check{\Pi}$ be a Metropolis–Hastings transition kernel with target $\check{\pi}$ and proposal distribution $\check{Q}_x$ where $x \in \check{\mathsf{X}}$. If Π is itself a Metropolis–Hastings kernel, it is possible to modify its proposal Q_x where $x \in \mathsf{X}$ in order to design $\check{Q}_x$. For example, one could choose, for some $w \in (0,1)$ and a probability measure $\mu : \mathcal{B}(\check{\mathsf{X}}) \to [0,1]$,

$$\check{Q}_x(\mathrm{d}y) = \mathbb{I}(x \in \mathsf{X})\{wQ_x(\mathrm{d}y) + (1-w)\delta_a(\mathrm{d}y)\} + \mathbb{I}(x = a)\mu(\mathrm{d}y). \tag{16}$$

The second strategy is to define, for some $w \in (0,1)$ and transition kernels Π_1 and Π_2,

$$\check{\Pi}(x, \mathrm{d}y) := w\Pi_1(x, \mathrm{d}y) + (1-w)\Pi_2(x, \mathrm{d}y),$$

where $\Pi_1(x, \mathrm{d}y) = \mathbb{I}(x \in \mathsf{X})\Pi(x, \mathrm{d}y) + \mathbb{I}(x = a)\delta_a(\mathrm{d}y)$ and Π_2 allows the chain to move between X and $\{a\}$. One choice of Π_2, suggested by Brockwell and Kadane (2005), is a Metropolis–Hastings kernel with proposal

$$Q_x(\mathrm{d}y) = \mathbb{I}(x \in \mathsf{X})\delta_a(\mathrm{d}y) + \mathbb{I}(x = a)\mu(\mathrm{d}y).$$

For both strategies, the choices of w and μ can greatly affect performance. Brockwell and Kadane (2005) suggest that adaptation of these parameters at regeneration times of $\check{\mathbf{Y}}$ can be beneficial, a procedure justified theoretically by Gilks et al. (1998). The strategies outlined here are not exhaustive. In Section 5 a different approach is taken, but with essentially the same idea of modifying Π in an appropriate way to define $\check{\Pi}$.

4.3 General applicability of the methodology

The perfect simulation algorithms we have developed rely on the uniform ergodicity of the Markov chain of interest. We show here that whenever $\mathbf{X}$ is uniformly ergodic, then $\check{\mathbf{X}}$ is uniformly ergodic when $\check{\Pi}$ is a modification of Π in a specific sense.

Proposition 5. *Assume that a generic Markov kernel* $\check{\Pi} : \check{\mathsf{X}} \times \mathcal{B}(\check{\mathsf{X}}) \to [0,1]$ *satisfies* $\check{\Pi}(a, \mathsf{X}) > 0$ *and for some* $w > 0$,

$$\check{\Pi}(x, A) \geq w\Pi(x, A), \quad x \in \mathsf{X}, A \in \mathcal{B}(\mathsf{X}). \tag{17}$$

Then $\mathbf{X}$ *being uniformly ergodic implies that* $\check{\mathbf{X}}$ *is uniformly ergodic but the converse does not hold.*

The conditions of Proposition 5 are met for both strategies in Section 4.2. This implies that one can modify the transition kernels of uniformly ergodic Markov chains on a general state space to construct transition kernels of uniformly ergodic Markov chains with an artificial singleton atom straightforwardly. Inspection of the proof of the converse failing to hold also suggests that in some cases one can even obtain perfect samples from the stationary distribution of a non-uniformly ergodic Markov chain by suitable definition of $\check{\Pi}$, but we do not pursue this further here.

The existence of a $\beta > 0$ such that $\underline{p} \geq \beta$ is guaranteed in general for uniformly ergodic Markov chains with a proper, accessible atom, as long as one is willing to consider a k-step transition kernel. Indeed, any such chain $\check{\mathbf{X}}$ with a proper, accessible atom $\{a\}$ and transition kernel $\check{\Pi}$ satisfies, for some $m \in \mathbb{N}$ and $d > 0$,

$$\inf_{x \in \check{\mathsf{X}}} \check{\Pi}^m(x, \{a\}) \geq d,$$

and $\lim_{k\to\infty} \inf_{x\in\check{\mathsf{X}}} \check{\Pi}^k(x, \{a\}) = \check{\pi}(\{a\})$. Therefore, one can increase $\underline{p}$ by choosing a large enough k and treating the Markov kernel $\check{\Pi}^k$ as the Markov kernel of interest. This presents no additional difficulties, since we only require the ability to simulate from $\check{\Pi}^k$ and identify the atom $\{a\}$. When

$\{a\}$ is an artificial atom, it may also be necessary to construct $\check{\Pi}$ in such a way that $\check{\pi}(\{a\})$ is not very close to 0 or 1. The existence of $\beta > 0$ does not of course imply that it is always known in practice but, at least in principle, one can simulate from the stationary distribution of any uniformly ergodic Markov chain $\mathbf{X}$ via the introduction of an artificial atom and the consideration of a k-step version of the modified transition kernel. In Section 5, a Markov kernel $\check{P}_N$ with an artificial atom is considered for which one can always treat $\check{P}_N$ itself as the Markov kernel of interest. Moreover, the subscript N is a parameter of this Markov kernel that controls its rate of convergence.

The uniform ergodicity requirement for application of Algorithms 4 or 5 might appear strong but a similar condition is required for the general CFTP algorithm presented in Propp and Wilson (1996), as established in (Foss and Tweedie, 1998, Theorem 4.2).

5 Perfect simulation according to a Feynman–Kac path measure

5.1 Feynman–Kac path measure

In this section, we focus on a generic discrete-time Feynman–Kac model with time horizon n. Let $(\mathsf{Z}, \mathcal{B}(\mathsf{Z}))$ be a measurable space. Consider a probability measure $\mu : \mathcal{B}(\mathsf{Z}) \to [0,1]$, some Markov kernels $M_p : \mathsf{Z} \times \mathcal{B}(\mathsf{Z}) \to [0,1]$ for $p \in \{2,\ldots,n\}$ and non-negative $\mathcal{B}(\mathsf{Z})$-measurable functions $G_p : \mathsf{Z} \to \mathbb{R}_+$ for $p \in \{1,\ldots,n\}$. We write $\mathcal{M} := (M_p)_{p\in\{2,\ldots,n\}}$ and $\mathcal{G} := (G_p)_{p\in\{1,\ldots,n\}}$. We define for any $p \in \{1,\ldots,n\}$, the measure γ_p by

$$\gamma_p(A) := \int_A \left[\prod_{q=1}^{p} G_q(z_q)\right] \mu(\mathrm{d}z_1) \prod_{q=2}^{p} M_q(z_{q-1}, \mathrm{d}z_q), \quad A \in \mathcal{B}(\mathsf{Z}^p), \tag{18}$$

and its associated probability measure $\pi_p := \gamma_p(1)^{-1}\gamma_p$. With $\mathsf{X} := \mathsf{Z}^n$ the Feynman–Kac path measure of interest is the probability measure $\pi := \pi_n$ on $\mathcal{B}(\mathsf{X})$.

Feynman–Kac models naturally accommodate HMMs (see, e.g., Del Moral, 2004). Indeed, consider a latent general state space Markov chain $(Z_p)_{p\geq 1}$ such that $Z_1 \sim \mu$ and $Z_p | Z_{p-1} \sim M_p(Z_{p-1}, \cdot)$ and observations $(Y_1,\ldots,Y_n) = (y_1,\ldots,y_n)$ where $(Y_1,\ldots,Y_n)$ are assumed conditionally independent given $(Z_1,\ldots,Z_n)$ with each Y_p depending on $(Z_1,\ldots,Z_n)$ only through Z_p with $Y_p | Z_p \sim g_p(Z_p, \cdot)$. If we let $G_p(z_p) = g_p(z_p, y_p)$, then π corresponds to the distribution of the states $(Z_1,\ldots,Z_n)$ of the unobserved Markov chain conditional upon the observations $(y_1,\ldots,y_n)$ and $\pi(f)$ is the associated conditional expectation of $f(Z_1,\ldots,Z_n)$.

More generally, Feynman–Kac models, through expressions such as (18), can be used to define arbitrary distributions π. The primary benefit of expressing distributions of interest in this way is that it allows the approximation of expectations, $\pi(f)$, using SMC methods. For example, a sophisticated SMC methodology for sampling from a complex distribution using a sequence of auxiliary "bridging" distributions and associated MCMC kernels is the SMC sampler methodology of Del Moral et al. (2006).

5.2 A particle filter

The general purpose SMC algorithm, often referred to as a particle filter, for estimating $\pi(f)$ was proposed in Stewart and McCarty Jr (1992), Gordon et al. (1993) and Kitagawa (1996) in the context of HMMs; see Doucet and Johansen (2011) for a recent survey. The algorithm is described in Algorithm 7, in which $\mathcal{C}(p_1,\ldots,p_N)$ generically denotes the categorical distribution over $\{1,\ldots,N\}$ with probabilities proportional to $(p_1,\ldots,p_N)$. We follow the presentation of Andrieu et al. (2010). For each $i \in \{1,\ldots,N\}$ and $p \in \{1,\ldots,n\}$, the variable ζ_p^i is the ith *particle* at time p; when $p > 1$, A_{p-1}^i is the index of the *ancestor* of this same particle, in the sense that in Algorithm 7,

$$\zeta_p^i \sim M_p(\zeta_{p-1}^{A_{p-1}^i}, \cdot).$$

This particular instance of SMC uses multinomial resampling, so-called because the ancestor indices $\{A_{p-1}^i : i \in \{1,\ldots,N\}\}$ are conditionally i.i.d. given $(\zeta_{p-1}^1,\ldots,\zeta_{p-1}^N)$ and distributed according $\mathcal{C}\left(G_{p-1}(\zeta_{p-1}^1),\ldots,G_{p-1}(\zeta_{p-1}^N)\right)$, implying that the "offspring" vector

$$\left(\sum_{i=1}^N \mathbb{I}(A_{p-1}^i = 1),\ldots,\sum_{i=1}^N \mathbb{I}(A_{p-1}^i = N)\right)$$

is multinomially distributed.

Algorithm 7 SMC: sampling $V \sim Q^N$

1. Simulate $\zeta_1^i \sim \mu$ for $i \in \{1,\ldots,N\}$.
2. [If $n > 1$] For $p = 2,\ldots,n$
   (a) Simulate $A_{p-1}^i \sim \mathcal{C}\left(G_{p-1}(\zeta_{p-1}^1),\ldots,G_{p-1}(\zeta_{p-1}^N)\right)$ for $i \in \{1,\ldots,N\}$.
   (b) Simulate $\zeta_p^i \sim M_p(\zeta_{p-1}^{A_{p-1}^i},\cdot)$ for $i \in \{1,\ldots,N\}$.
3. Set $V = (\zeta_1^1,\ldots,\zeta_n^N,A_1^1,\ldots,A_{n-1}^N)$.

In order to define an estimate of the path measure π as a function of the random variables produced in Algorithm 7, it is necessary to define the *ancestral lineage* of a particle index. Following Andrieu et al. (2010), for any $k \in \{1,\ldots,N\}$, we define B^k to be the $\{1,\ldots,N\}^n$-valued random variable satisfying $B_n^k := k$ and $B_p^k := A_p^{B_{p+1}^k}$. The random variable

$$\zeta^k := (\zeta_1^{B_1^k},\ldots,\zeta_n^{B_n^k}) \tag{19}$$

is then a path taking values in X. With this notation, a natural estimate of π is

$$\pi^N(\mathrm{d}x) := \frac{\sum_{k=1}^N G_n(\zeta_n^k)\delta_{\zeta^k}(\mathrm{d}x)}{\sum_{j=1}^N G_n(\zeta_n^j)}, \tag{20}$$

and we estimate accordingly $\pi(f)$ using $\pi^N(f)$. Many asymptotic properties of $\pi^N(f)$ have been established in the literature (see Del Moral, 2004, for a comprehensive treatment). For example, under very mild conditions $\pi^N(f)$ converges almost surely to $\pi(f)$ and a $\sqrt{N}$-central limit theorem holds for $\pi^N(f)$ as $N \to \infty$ (Del Moral and Miclo, 2000; Chopin, 2004).

5.3 Motivation for perfect simulation

The convergence results mentioned above establish that (20) can be used to approximate the probability measure π. In fact, it is clear that if $\pi^N(A)$ is "close" to $\pi(A)$ for any $A \in \mathcal{B}(\mathsf{X})$, then the marginal distribution of a single path ζ^K where $K \sim \mathcal{C}\left(G_n(\zeta_n^1),\ldots,G_n(\zeta_n^N)\right)$ is "close" to π. We will see that under mild conditions, this "closeness" can be controlled by choosing a large enough value of N.

Letting $[N] := \{1,\ldots,N\}$, we define $\mathsf{V}_N := \mathsf{Z}^{Nn} \times [N]^{N(n-1)}$ and $Q^N : \mathcal{B}(\mathsf{V}_N) \to [0,1]$ to be the probability measure associated with the variables produced by running the SMC scheme with N particles described in Algorithm 7. We then define a collection of conditional probability measures $\{Q_v^N : v \in \mathsf{V}_N\}$ such that $Q_v^N : \mathcal{B}([N]^n \times \mathsf{X}) \to [0,1]$ is the probability measure associated with picking an ancestral lineage and a path. We show how to sample from this measure in Algorithm 8. A joint probability measure $Q^N : \mathcal{B}([N]^n \times \mathsf{X} \times \mathsf{V}_N) \to [0,1]$ is then defined as $Q^N(\mathbf{k},\mathrm{d}x,\mathrm{d}v) := Q^N(\mathrm{d}v)Q_v^N(\mathbf{k},\mathrm{d}x)$ and the marginal distribution of X when $(X,V) \sim Q^N$ is also denoted by $Q^N : \mathcal{B}(\mathsf{X}) \to [0,1]$ with $Q^N(\mathrm{d}x) := \sum_{\mathbf{k}\in[N]^n} \int_{\mathsf{V}_N} Q^N(\mathbf{k},\mathrm{d}x,\mathrm{d}v)$.

Algorithm 8 Pick a path: sampling $(\mathbf{K}, X) \sim Q_v^N$

Letting $v = (z_1^1, \ldots, z_n^N, a_1^1, \ldots, a_{n-1}^N)$,

1. Sample $K_n \sim \mathcal{C}\left(G_n(z_n^1), \ldots, G_n(z_n^N)\right)$ and [if $n > 1$] for $p = n-1, n-2, \ldots, 1$ set $K_p = a_p^{K_{p+1}}$.
2. Set $\mathbf{K} = (K_1, \ldots, K_n)$ and $X = \left(z_1^{K_1}, \ldots, z_n^{K_n}\right)$.

In order to show that Q^N is indeed "close" to π, and because the construction itself is a central component of the π-invariant Markov kernel we shall be concerned with, we introduce a collection of probability measures $\{\bar{Q}_x^N : x \in \mathsf{X}\}$. For a given $x \in \mathsf{X}$, called the reference path, $\bar{Q}_x^N : \mathcal{B}([N]^n \times \mathsf{V}_N) \to [0,1]$ corresponds to the distribution of the ancestral lineage of the reference path and the V_N-valued variables produced in a *conditional* SMC algorithm, a procedure introduced in Andrieu et al. (2010) and described in Algorithm 9 that has both theoretical and methodological applications. We will also denote by $\bar{Q}_x^N : \mathcal{B}(\mathsf{V}_N) \to [0,1]$ the probability measure associated with the marginal distribution of the V_N-valued variables produced, i.e. $\bar{Q}_x^N(\mathrm{d}v) := \sum_{\mathbf{k} \in [N]^n} \bar{Q}_x^N(\mathbf{k}, \mathrm{d}v)$.

Algorithm 9 Conditional SMC: sampling $(\mathbf{K}, V) \sim \bar{Q}_x^N$

1. Sample K_p independently and uniformly on $\{1, \ldots, N\}$ for each $p \in \{1, \ldots, n\}$.
2. Set $\zeta_1^{K_1} = x_1$ and simulate $\zeta_1^i \sim \mu$ for $i \in \{1, \ldots, N\} \setminus \{K_1\}$.
3. [If $n \geq 1$] For $p = 2, \ldots, n$
   (a) Set $A_{p-1}^{K_p} = K_{p-1}$ and simulate $A_{p-1}^i \sim \mathcal{C}\left(G_{p-1}(\zeta_{p-1}^1), \ldots, G_{p-1}(\zeta_{p-1}^N)\right)$ for $i \in \{1, \ldots, N\} \setminus \{K_p\}$.
   (b) Set $\zeta_p^{K_p} = x_p$ and simulate $\zeta_p^i \sim M_p(\zeta_{p-1}^{A_{p-1}^i}, \cdot)$ for $i \in \{1, \ldots, N\} \setminus \{K_p\}$.
4. Set $\mathbf{K} = (K_1, \ldots, K_n)$ and $V = (\zeta_1^1, \ldots, \zeta_n^N, A_1^1, \ldots, A_{n-1}^N)$.

Central to some of the methodology in Andrieu et al. (2010) is the behaviour of a random variable known as the normalizing constant estimate, renormalized so that it has expectation 1 when evaluated at $(\mathbf{K}, X, V) \sim Q^N$ (Andrieu and Roberts, 2009; Andrieu and Vihola, 2012; Lee and Łatuszyński, 2012; Andrieu et al., 2013). Letting $v = (z_1^1, \ldots, z_n^N, a_1^1, \ldots, a_{n-1}^N)$, this renormalized estimate is the Radon–Nikodym derivative

$$\phi^N(v) := \frac{\pi(\mathrm{d}x)\bar{Q}_x^N(\mathbf{k}, \mathrm{d}v)}{Q^N(\mathbf{k}, \mathrm{d}x, \mathrm{d}v)} = \frac{1}{\gamma_n(1)} \prod_{p=1}^{n} \frac{1}{N} \sum_{j=1}^{N} G_p(z_p^j), \tag{21}$$

which depends neither on the value of the picked particle X nor on its ancestral lineage $\mathbf{K}$. Letting E^N denote expectation w.r.t. the law of Q^N and $\bar{\mathsf{E}}_x^N$ denote expectation w.r.t. the law of $\bar{Q}_x^N$, the following result provides a non-asymptotic upper bound on the Radon–Nikodym derivative between π and Q^N.

Proposition 6. *Assume that for each $p \in \{1, \ldots, n\}$, $G_p(z_p) > 0$ for all $z_p \in \mathsf{Z}$. Then*

$$\frac{\pi(\mathrm{d}x)}{Q^N(\mathrm{d}x)} \leq \int_{\mathsf{V}_N} \phi^N(v)\bar{Q}_x^N(\mathrm{d}v) = \bar{\mathsf{E}}_x^N\left[\phi^N(V)\right].$$

Remark. A similar, but slightly more complicated expression holds without the strict positivity of the potentials.

Proposition 6, together with results in Andrieu et al. (2013, Section 4) implies a uniform upper bound on $\pi(\mathrm{d}x)/Q^N(\mathrm{d}x)$ when the potentials in $\mathcal{G}$ are additionally uniformly bounded above. In

practice, SMC methods work well for models which have "forgetting properties". This notion can be quantified, and used to bound quantities that arise in their analysis. In particular we define for each $p \in \{1, \ldots, n\}$ and $k \in \{0, \ldots n-p\}$ the function $F_{p,k} : \mathsf{Z} \to \mathbb{R}_+$ by

$$F_{p,k}(z_p) := \frac{\gamma_{p-1}(1)}{\gamma_{p+k}(1)} \int_{\mathsf{Z}^k} \left[\prod_{q=p}^{p+k} G_q(z_q)\right] \left[\prod_{q=p+1}^{p+k} M_q(z_{q-1}, \mathrm{d}z_q)\right],$$

with the convention that $\gamma_0(1) = 1$. The value $F_{p,k}(z_p)$ can be viewed as an appropriately normalized expectation of the product of the $k+1$ potentials evaluated along a path started at $Z_p = z_p$ and evolving according to the dynamics in $\mathcal{M}$ for k steps. In the context of an HMM, e.g., $F_{p,k}(z_p)$ is the likelihood of the $k+1$ observations $(y_p, \ldots, y_{p+k})$ conditional upon $Z_p = z_p$ divided by the likelihood of the same observations conditional upon $y_1, \ldots, y_{p-1}$, and indeed one can see that if $\sup_{z_p \in \mathsf{Z}} F_{p,k}(z_p)$ eventually stabilizes as k increases then the HMM "forgets" in a particular sense the distribution of Z_p at time $p+k$. The level of forgetting of the finite time Feynman–Kac model as a whole can be crudely, but succinctly summarized in the single quantity

$$F := \max_{p \in \{1,\ldots,n\}, k \in \{0,\ldots,n-p\}} \sup_{z_p \in \mathsf{Z}} F_{p,k}(z_p),$$

which is necessarily finite whenever the potentials in $\mathcal{G}$ are upper bounded.

Proposition 7. *Assume there exists $B < \infty$ such that for each $p \in \{1, \ldots, n\}$, $0 < G_p(z_p) < B$ for all $z_p \in \mathsf{Z}$. Then for any $N \geq 2$,*

$$\sup_{x \in \mathsf{X}} \frac{\pi(\mathrm{d}x)}{Q^N(\mathrm{d}x)} \leq \bar{\mathsf{E}}_x^N \left[\phi^N(V)\right] \leq \left(1 + \frac{2(F-1)}{N}\right)^n. \tag{22}$$

The second inequality in (22) is tight. In interpreting (22), it is informative to consider the effect of doubling the number of particles N. Letting $C = 2(F-1)$, we have $\left(1 + \frac{C}{N}\right)^n / \left(1 + \frac{C}{2N}\right)^n = \left(\frac{2N+2C}{2N+C}\right)^n = \left(1 + \frac{C}{2N+C}\right)^n$ which is approximately 2^n if $C \gg 2N$ and $1 + \frac{nC}{2N}$ when $2N \gg C$. Hence the improvement is close to exponential for small N but asymptotically linear.

While in general F will grow exponentially with n when the potentials have a uniform upper bound, under stronger assumptions it can be shown that F has an upper bound that is independent of n. In practice, the performance of these algorithms in a large class of statistical applications suggests that $\bar{\mathsf{E}}_x^N \left[\phi^N(V)\right]$ grows as a small power of n, and there is some recent theory pointing towards this direction (Lindsten et al., 2014, Section 4.3). One can think of F as providing a rough indication of the level of difficulty of the Feynman-Kac model, in terms of the closeness of Q^N to π, and more generally as quantifying the suitability of SMC methodology for a given application.

The form of (22) immediately implies that rejection sampling to obtain perfect samples according to π is possible *in principle*. However, since it is not possible to evaluate $\pi(\mathrm{d}x)/Q^N(\mathrm{d}x)$ exactly such a procedure has no existing practical implementation. A natural way to bypass the calculation of $\pi(\mathrm{d}x)/Q^N(\mathrm{d}x)$, consists of using rejection sampling to sample from the joint distribution defined by $\pi(\mathrm{d}x)\bar{Q}_x^N(\mathrm{d}v)$, with proposal $Q^N(\mathrm{d}x, \mathrm{d}v)$. One can accept proposals with probability $M_N^{-1}\phi^N(V)$, where $M_N \geq \sup_v \phi^N(v)$. From (21) we can see that $M = \gamma_n(1)^{-1} \prod_{p=1}^n \sup_x G_p(x)$, which is independent of N and typically grows exponentially in n, is usually the smallest value satisfying this condition. The expected acceptance probability $\mathsf{E}^N \left[M^{-1}\phi^N(V)\right] = M^{-1}$ is unfortunately independent of N and so this algorithm appears unable to take advantage of (22).

The combination of the potential of Algorithm 7 to produce perfect samples and the inability to evaluate $\pi(\mathrm{d}x)/Q^N(\mathrm{d}x)$ invites modifications of Algorithm 7. In Andrieu et al. (2012), a perfect sampling method is proposed where the mechanism governing particle offspring is fundamentally changed from selection with a constant population size at each time to stochastic branching; their method is able to produce perfect samples from some models (see also Section 6.6). We will instead use a π-invariant Markov kernel proposed in Andrieu et al. (2010) whose main ingredient is Algorithm 9 and is, in many ways, closer in spirit to Algorithm 7.

5.4 Iterated conditional SMC kernel

The iterated conditional SMC kernel (Andrieu et al., 2010) is a Markov kernel $P_N : \mathsf{X} \times \mathcal{B}(\mathsf{X}) \to [0,1]$ defined by

$$P_N(x, \mathrm{d}y) := \int_{\mathsf{V}_N} \bar{Q}_x^N(\mathrm{d}v) Q_v^N(\mathrm{d}y).$$

That is, we first run a conditional SMC algorithm (Algorithm 9) to produce the auxiliary variables V and then choose a path (Algorithm 8) given the auxiliary variables V just produced. This Markov kernel leaves π invariant, see, e.g., Andrieu et al. (2013, Lemma 4).

Our interest in using this particular Markov kernel to facilitate regenerations and perfect simulation stems from both its good empirical performance and its uniform ergodicity properties. In particular, Chopin and Singh (2013) provided the first such result implying that under some assumptions $P_N(x, \mathrm{d}y) \geq \epsilon_N \pi(\mathrm{d}y)$ where $\epsilon_N \to 1$ but without a rate of convergence. Andrieu et al. (2013) and Lindsten et al. (2014) provide convergence under the weaker assumption that the potentials in $\mathcal{G}$ are π-essentially upper bounded with quantitative rates of convergence for $\epsilon_N \to 1$. In addition, Andrieu et al. (2013) showed that this essential boundedness condition is also necessary for P_N to define a uniformly ergodic Markov chain.

Andrieu et al. (2013, Corollary 15) provides the following bound on ϵ_N in terms of F

$$P_N(x, \mathrm{d}y) \geq \left(\frac{N-1}{N+2(F-1)}\right)^n \pi(\mathrm{d}y) = \left(1 - \frac{2F-1}{N+2(F-1)}\right)^n \pi(\mathrm{d}y), \tag{23}$$

which implies that $\epsilon_N > 0$ whenever F is finite and $N \geq 2$. This bound can be compared with (22), and behaves asymptotically in N as $(1 - \frac{(2F-1)n}{N}) + \mathcal{O}(N^{-2})$ but for large n and large F the improvement by increasing N is drastic. Finally, under fairly strong assumptions on the elements of μ, $\mathcal{M}$ and $\mathcal{G}$, F does not grow with n and it is possible to control ϵ_N as $n \to \infty$ by scaling N only linearly with n. In particular, if $N = Bn$ then $\epsilon_N \geq \exp\left(-\frac{2(F-1)}{B}\right)$ (Andrieu et al., 2013, Corollary 15). In order to guarantee $\epsilon_N \geq 0.75$, e.g., we can take $B = -2(F-1)/\log(0.75)$, which is slightly less than $7(F-1)$ and so ensuring that ϵ_N is fairly large requires a computational complexity in $\mathcal{O}(n^2)$.

5.5 Atomic extensions of the Feynman–Kac path measure and iterated conditional SMC kernel

Let $\check{\mathsf{Z}} := \mathsf{Z} \cup \alpha$, where $\alpha = \{a\}$ and a is a distinguished point, and $\check{\mathsf{X}} := \check{\mathsf{Z}}^n$, $a_n := (a, \ldots, a)$. We propose here a generic way to define a new probability measure $\check{\pi}$ on $\check{\mathsf{X}}$ which satisfies for some $k \in (0,1)$,

$$\check{\pi}(A) = k\pi(A) + (1-k)\mathbb{I}(a_n \in A), \quad A \in \mathcal{B}(\check{\mathsf{X}}), \tag{24}$$

where π is the Feynman–Kac path measure. We do not follow the approach suggested in Section 4 to build $\check{\pi}$ but directly modify the original Feynman–Kac model to introduce the atom $\alpha_n := \{a_n\}$. The benefit of this approach is that it allows us to come up straightforwardly with a $\check{\pi}$-invariant Markov kernel with well-understood properties.

The extended Feynman–Kac model is defined by the initial distribution $\check{\mu}$, the Markov kernels $\check{\mathcal{M}} := (\check{M}_p)_{p \in \{2,\ldots,n\}}$ and potential functions $\check{\mathcal{G}} := (\check{G}_p)_{p \in \{1,\ldots,n\}}$ on $\check{\mathsf{Z}}$ which are given by

$$\begin{aligned}
\check{\mu}(A) &:= (1-b)\mu(A \cap \mathsf{Z}) + b\mathbb{I}\{a \in A\}, \\
\check{M}_p(x, A) &:= M_p(x, A)\mathbb{I}\{x \in \mathsf{X}\} + \mathbb{I}\{x = a, A = \alpha\}, \\
\check{G}_p(x) &:= G_p(x)\mathbb{I}\{x \in \mathsf{X}\} + \psi_p \mathbb{I}\{x = a\},
\end{aligned}$$

for $A \in \mathcal{B}(\check{\mathsf{Z}})$ where $b \in (0,1)$ and $\psi_1, \ldots, \psi_n$ are user-defined positive constants. Similar to the definition of γ_n in Section 5.1, we define the measure $\check{\gamma}_n$ by

$$\check{\gamma}_n(A) := \int_A \left[\prod_{p=1}^n \check{G}_p(x_p)\right] \check{\mu}(\mathrm{d}x_1) \prod_{p=2}^n \check{M}_p(x_{p-1}, \mathrm{d}x_p), \quad A \in \mathcal{B}(\check{\mathsf{X}}),$$

and its associated probability measure $\check{\pi} := \check{\gamma}_n(1)^{-1}\check{\gamma}_n$. It follows easily from these definitions that (24) holds with

$$k = \frac{1-b}{1-b+b\gamma_n(1)^{-1}\prod_{p=1}^n \psi_p}. \tag{25}$$

The transition kernel we propose to use to sample from the extended Feynman–Kac path measure $\check{\pi}$ is simply the i-cSMC kernel targetting $\check{\pi}$ instead of π. We denote by $\check{P}_N$ this kernel, where N is the number of particles used. The uniform ergodicity of the Markov chain with transition kernel $\check{P}_N$ is straightforward, and indeed the Markov chain defined by $\check{P}_N$ is uniformly ergodic if and only if the Markov chain defined by P_N is uniformly ergodic. The forgetting properties of the extended Feynman–Kac model can be quantified via

$$\check{F}_{p,k}(z_p) := \frac{\check{\gamma}_{p-1}(1)}{\check{\gamma}_{p+k}(1)} \int_{\check{\mathsf{Z}}^k} \left[\prod_{q=p}^{p+k} \check{G}_q(z_q)\right] \left[\prod_{q=p+1}^{p+k} \check{M}_q(z_{q-1}, \mathrm{d}z_q)\right],$$

and

$$\check{F} := \max_{p\in\{1,\ldots,n\}, k\in\{0,\ldots,n-p\}} \sup_{z_p\in\mathsf{Z}} \check{F}_{p,k}(z_p).$$

At this point, it is clear that the notion of forgetting encoded by F and $\check{F}$ is not the same as forgetting of the initial distribution of a Feynman–Kac model, which refers to the time n marginal of π_n being asymptotically independent of μ as $n \to \infty$ under suitable assumptions. Indeed, one can see that conditional upon $Z_p = a$, the distribution of Z_{p+k} is a point mass at a, regardless of k, and yet we can still have "forgetting" in this weaker sense.

Proposition 8. *The Markov chain with transition kernel $\check{P}_N$ is uniformly ergodic if and only if the potentials in $\mathcal{G}$ are π-essentially bounded and $N \geq 2$. Moreover, for any $N \geq 2$,*

$$\check{P}_N(x, A) \geq \check{\epsilon}_N \check{\pi}(A), \quad x \in \check{\mathsf{X}}, A \in \mathcal{B}(\check{\mathsf{X}}), \tag{26}$$

where $\check{\epsilon}_N > 0$ and $\check{\epsilon}_N \to 1$ as $N \to \infty$. One can take $\check{\epsilon}_N = \left(\frac{N-1}{N+2(F-1)}\right)^n$.

In order to implement either Algorithm 2 or Algorithm 3, one needs a lower bound on $\inf_{x\in\check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$. This can be obtained via a lower bound on the quantities $\check{\pi}(\alpha_n)$ and a value of $\check{\epsilon}_N$ satisfying (26). In general, it may not be possible to obtain such bounds analytically using current theory, although Proposition 8 is reassuring in that one knows that $\check{\epsilon}_N$ can be made arbitrarily close to 1 by increasing N. We will see in the next section that a specific choice of the constants $\psi_1, \ldots, \psi_n$ ensures that $\check{F}$ is equal to F, and while this choice is typically not possible one can select $\psi_1, \ldots, \psi_n$ stochastically so that they are close to this specific choice.

5.6 Parameter settings

As mentioned earlier in Section 4, it is desirable that $\check{\pi}(\alpha_n)$ is neither too close to 0 nor too close to 1. We outline here a simple strategy for selecting b and $\psi_1, \ldots, \psi_n$ to achieve this. We note that if $\prod_{p=1}^n \psi_p = \gamma_n(1)$ then we obtain $\check{\gamma}_n(1) = \gamma_n(1)$ and $k = 1 - b$ in (25). We therefore suggest choosing $b = \frac{1}{2}$ with the constants $\psi_1, \ldots, \psi_n$ chosen such that $\prod_{p=1}^n \psi_p \approx \gamma_n(1)$. The behaviour of the i-cSMC Markov chain associated with the extended model specified by $\check{\mu}$, $\check{\mathcal{M}}$ and $\check{\mathcal{G}}$ can be related to the behaviour of the i-cSMC Markov chain associated with the model specified by μ, $\mathcal{M}$ and $\mathcal{G}$. In particular, we can relate the forgetting properties of both models by relating the constants F and $\check{F}$ associated with the two models.

Proposition 9. *Assume that for some $E \geq 1$*

$$\max_{p\in\{1,\ldots,n\}} \max\left\{\frac{\psi_p}{\gamma_p(1)/\gamma_{p-1}(1)}, \frac{\gamma_p(1)/\gamma_{p-1}(1)}{\psi_p}\right\} \leq E.$$

Then $\check{F} \leq FE^n$.

Although the dependence of this bound on n is exponential, for models in which F is not too large, small values of E result with high probability when each ψ_p is an SMC estimate of $\gamma_p(1)/\gamma_{p-1}(1)$. Practically, we run Algorithm 7 with a large number N' of particles for the model specified by μ, $\mathcal{M}$ and $\mathcal{G}$ and approximate each $\gamma_p(1)/\gamma_{p-1}(1)$ by $\psi_p := \frac{1}{N'}\sum_{i=1}^{N'} G_p(\zeta_p^i)$, which is a consistent estimate as $N' \to \infty$. Then one can run Algorithm 7 with N' particles for the model specified by $\check{\mu}$, $\check{\mathcal{M}}$ and $\check{\mathcal{G}}$ and estimate $\check{\pi}(\alpha_n)$ by $\frac{1}{N'}\sum_{i=1}^{N'} \mathbb{I}(\zeta_n^i = a)$. Since these are only estimates, it is prudent to conservatively use a value smaller than this latter estimate, e.g., by using a concentration inequality after multiple runs of Algorithm 7 for the model specified by $\check{\mu}$, $\check{\mathcal{M}}$ and $\check{\mathcal{G}}$. Ideally, this estimate is close to b, as suggested by (25). In practice, when $\check{\pi}(\alpha_n)$ is close to b, this is also a good indication that each ψ_p is very close to $\gamma_p(1)/\gamma_{p-1}(1)$ since

$$\gamma_n(1) = \prod_{p=1}^{n} \frac{\gamma_p(1)}{\gamma_{p-1}(1)} \approx \prod_{p=1}^{n} \psi_p.$$

Approximating or bounding F is challenging, but explicit upper bounds can be obtained under strong assumptions (see Appendix B). Although the finiteness of F is guaranteed when n is finite whenever the potentials in $\mathcal{G}$ are bounded, little attention has been paid in the literature to obtaining explicit upper bounds that are reasonably small. For example, Whiteley (2013) provides only the existence of $F < \infty$ as $n \to \infty$ but no explicit bounds. One situation in which an upper bound on F can be sufficiently small to be of practical use is discussed in Section 7.

In practice, we have found that this procedure for selecting $\psi_1, \ldots, \psi_n$ and estimating bounds on $\check{\epsilon}_N$ and $\check{\pi}(\alpha_n)$ is very effective, even though deterministic bounds on $\check{\epsilon}_N$ and $\check{\pi}(\alpha_n)$ are not available. Since this procedure does not provide a lower bound on $\inf_{x\in\check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$, one can only hope for a bound that holds with high probability. In Sections 6.1 and 6.2 we discuss diagnostics and estimation of $\inf_{x\in\check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$, and the sensitivity of the procedure to choosing a value of β superior to $\inf_{x\in\check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$.

6 Discussion

6.1 Diagnostics and estimation of β

In some situations, it may not be possible to obtain a lower bound β on $\underline{p}$. One can nevertheless attempt to estimate such a bound and there are many strategies one may use to perform this estimation. In general, one seeks to find the maximum of the function $g : \mathsf{X} \to [0, 1]$ defined by $x \mapsto 1 - p(x)$, and one can obtain unbiased estimates of $g(x)$ for any $x \in \mathsf{X}$ by sampling from $\Pi(x, \cdot)$. This suggests that a stochastic optimization procedure could be used to estimate $\underline{p}$. An alternative approach is to simulate the Markov chain $\mathbf{X}$ for a long period of time and use the chain to estimate $\underline{p}$. It is possible then to use the same realization of the Markov chain to impute the regeneration indicators, i.e., run Algorithm 4 retrospectively.

It is also possible to monitor the validity of the assumption that a chosen β satisfies $\beta \leq \underline{p}$ using a simple and fairly inexpensive diagnostic. At each state x visited during the course of Algorithm 4 or 5, one can simply simulate $p(x)$-coins until their average exceeds β. The following result shows that this strategy will require little computational effort if $\underline{p} > \beta$ but that a failure to terminate after a reasonable number of $p(x)$-coins have been simulated indicates that $\underline{p}$ may be inferior to β.

Proposition 10. *Let $p \in [0, 1]$ and $(B_i)_{i\geq 1}$ be i.i.d.* Bernoulli(p) *random variables. Let τ be the stopping time defined by $\tau := \inf\{n \geq 1 : \frac{1}{n}\sum_{i=1}^{n} B_i > \beta\}$. Then the following results hold:*

1. *If $p < \beta$, then $\mathbb{P}(\tau < \infty) < 1$.*
2. *If $p = \beta$, then $\mathbb{P}(\tau < \infty) = 1$ and $\mathbb{E}(\tau) = \infty$.*
3. *If $p > \beta$, then $\mathbb{P}(\tau < \infty) = 1$ and $\mathbb{E}(\tau) \leq (1 - \beta)/(p - \beta)$.*

These properties immediately imply that repeated use of Algorithms 4 or 5 in tandem with running the diagnostic will not cause problems if $\beta < \underline{p}$ but that if $\beta = \pi - \operatorname{ess\,inf}_{x\in\mathsf{X}} p(x)$ each call has an infinite expected running time and if $\beta > \pi - \operatorname{ess\,inf}_{x\in\mathsf{X}} p(x)$ then with probability 1, one of the calls will eventually fail to terminate.

Corollary 2. *If the diagnostic is performed at every state visited by Algorithm 4 or 5, then:*

1. *If $\pi - \operatorname{ess\,inf}_{x\in\mathsf{X}} p(x) < \beta$, then the algorithm will not terminate with positive probability.*
2. *If $\pi - \operatorname{ess\,inf}_{x\in\mathsf{X}} p(x) = \beta$, then the algorithm has infinite expected running time.*
3. *If $\pi - \operatorname{ess\,inf}_{x\in\mathsf{X}} p(x) > \beta$, then the algorithm has finite expected running time.*

In the special case where $\beta = 1/m$ for some $m \in \mathbb{N}$, one can express analytically $\mathbb{P}(\tau < \infty)$. The following result shows that when p is even slightly different to β then the probability of not stopping in finite time is quite large. For example, if $\beta = 0.2$ and $p = 0.19$ then the probability of not stopping is over 0.06. We note that one can also determine probabilities associated with hitting times of certain subsets of $[0, 1]$ in this specific setting, but do not pursue this further.

Proposition 11. *Let $p \in [0, \beta)$ where $\beta = m^{-1}$ for some $m \in \mathbb{N}$, $(B_i)_{i\geq 1}$ be i.i.d.* Bernoulli(p) *random variables and $\tau := \inf\{n \geq 1 : \frac{1}{n}\sum_{i=1}^{n} B_i > \beta\}$. Then $\mathbb{P}(\tau < \infty) = p(m-1)/(1-p)$.*

Finally, in Section 3.4 it was seen that the computational complexity of the algorithms in Section 3.2 is improved when β is not too close to 0. In general, therefore, it may be the case that $\underline{p}$ is too small for the algorithms to be practical, even if $\mathbf{X}$ is uniformly ergodic. Following the remarks in Section 4.3, one general strategy is to treat the Markov kernel Π^k as the Markov kernel of interest. This can also be seen as an alternative to increasing N in the iterated conditional SMC kernel of Section 5.4, and may be appropriate when N is already sufficiently large that ϵ_N in (26) is reasonably large.

6.2 Sensitivity to the choice of ϵ

Consider the general case where (1) holds for $\nu = \delta_a$ and $\underline{p} = \underline{\epsilon} > 0$ but one has attempted to run Algorithm 3 with $s = \epsilon > \underline{\epsilon}$. We investigate here how this affects the samples obtained. It is clear that $R_{a,\epsilon}$ is not a Markov kernel but one can instead view the algorithm as sampling from

$$\tilde{R}_{a,\epsilon}(x, \mathrm{d}y) := \frac{\Pi(x, \mathrm{d}y) - \min\{\epsilon, p(x)\}\delta_a(\mathrm{d}y)}{1 - \min\{\epsilon, p(x)\}},$$

which results from flipping a $\min\{1, [1-p(x)]/(1-\epsilon)\}$-coin when deciding which mixture component in (12) to sample from. The corresponding Markov transition kernel whose invariant distribution we obtain perfect samples from is

$$\tilde{\Pi}(x, \mathrm{d}y) = \epsilon\delta_a(\mathrm{d}y) + (1-\epsilon)\tilde{R}_{a,\epsilon}(x, \mathrm{d}y),$$

which defines a uniformly ergodic Markov chain of invariant distribution denoted $\tilde{\pi}$. The following result indicates that closeness of ϵ and $\underline{\epsilon}$ implies closeness of π and $\tilde{\pi}$.

Proposition 12. *The invariant distributions π and $\tilde{\pi}$ satisfy*

$$\|\tilde{\pi} - \pi\|_{\mathrm{TV}} \leq 1 - \underline{\epsilon}/\epsilon,$$

where for a signed measure μ on $\mathcal{B}(\mathsf{X})$, $\|\mu\|_{\mathrm{TV}} := \sup_{A\in\mathcal{B}(\mathsf{X})} |\mu(A)|$.

The approximate Markov kernel $\tilde{R}_{a,\epsilon}(x, \cdot)$ is not exactly what is sampled from in Algorithm 5 when $p(x) < \epsilon$, as the latter is dependent on the specific implementation of the Bernoulli factory of a $(1-p)/(1-\epsilon)$-coin. However, Proposition 12 provides a reasonable approximation of the discrepancy between π and the distribution of the samples obtained from Algorithm 5 when $\underline{p} < \epsilon$.

6.3 Unbiased estimation of $\pi(f)$ with a perfect sample from π

With a perfect sample $X_1 \sim \pi$, it is trivial that $f(X_1)$ is an unbiased estimate of $\pi(f)$, where $f : \mathsf{X} \to \mathbb{R}$. To obtain a consistent estimate of $\pi(f)$, one can simply run any irreducible, π-invariant Markov kernel P initialized at X_1 and compute the unbiased MCMC estimate $n^{-1}\sum_{i=1}^{n} f(X_i)$. It may be less obvious that one can also combine a perfect sample with a number of imperfect samples to obtain an unbiased and consistent estimate of $\pi(f)$. Consider the π-invariant, i-cSMC kernel of Section 5.4. This kernel can be expressed as

$$P_N(x, \mathrm{d}y) := \int_{\mathsf{X}^N} \tilde{P}_N(x, \mathrm{d}y_{1:N}) \frac{\sum_{k=1}^{N} G(y_k)\delta_{y_k}(\mathrm{d}y)}{\sum_{j=1}^{N} G(y_j)},$$

where $\tilde{P}_N(x, \mathrm{d}y_{1:N}) := \int_{\mathsf{V}_N} \bar{Q}_x^N(\mathrm{d}v) \prod_{k=1}^{N} \delta_{\zeta^k}(\mathrm{d}y_k)$, with ζ^k defined in (19), and for $z \in \mathsf{X} = \mathsf{Z}^n$, $G(z) = G_n(z_n)$. It is clear that P_N satisfies $\int_{\mathsf{X}} \pi(\mathrm{d}x)P_N(x, \mathrm{d}z) = \pi(\mathrm{d}z)$, that $\int_{\mathsf{X}^2} f(y)\pi(\mathrm{d}x)P_N(x, \mathrm{d}y) = \pi(f)$ and, therefore, that

$$\int_{\mathsf{X}^{N+1}} \pi(\mathrm{d}x)\tilde{P}_N(x, \mathrm{d}y_{1:N}) \frac{\sum_{k=1}^{N} G(y_k)f(y_k)}{\sum_{j=1}^{N} G(y_j)} = \pi(f).$$

Hence, sampling $(Y_1, \ldots, Y_N) \sim \tilde{P}_N(X_1, \cdot)$ implies that the "self-normalized" estimate

$$\frac{\sum_{k=1}^{N} G(Y_k)f(Y_k)}{\sum_{j=1}^{N} G(Y_j)} \tag{27}$$

is an unbiased estimate of $\pi(f)$. The estimate (27) can be viewed as a specific case of a type of estimate studied in Whiteley and Lee (2014, Section 3) (which are in general, however, not unbiased), and a $\sqrt{N}$-central limit theorem holds for (27) under fairly mild assumptions. Moreover, the asymptotic variance is identical to that associated with $\pi^N(f)$, the standard particle filter estimate in Section 5.2 (Whiteley and Lee, 2014, Remark 7).

6.4 Non-uniformly ergodic Π and unbiased estimation of $\pi(f)$

In cases where $\underline{p} = 0$ or Π is not even uniformly ergodic, the more generally applicable strategy of Algorithm 1 can still be used. Implementations of this algorithm have typically sampled from $\eta_n^{(\nu,s)}$ in (7) by rejection using $\nu R_{\nu,s}^{n-1}$ as a proposal, which leads to an overall expected infinite running time.

An alternative procedure, however, is apparent from viewing $\eta_n^{(\nu,s)}$ as a marginal distribution associated with a Feynman–Kac path measure. That is, one can in principle sample from $\eta_n^{(\nu,s)}$ using either of the algorithms in Section 3.2 by constructing, e.g., an i-cSMC kernel $P_N^{(n)}$ with an artificial atom. In Section 5.4 it was seen that the computational time required to obtain such a sample could be in $\mathcal{O}(n^2)$, at least under strong assumptions. More generally, we note that a finite expected computational time of Algorithm 1 is guaranteed for a large class of Markov chains as long as one can sample N according to (8) in finite expected time and from $\eta_n^{(\nu,s)}$ in expected time polynomial in n.

Proposition 13. *Assume $\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)/\mathbb{E}_{\nu,s}(\tau_{\nu,s}) = \mathcal{O}(n^{-b})$ for some $b > 0$ and that one can sample according to (8) in finite expected time. Then if the expected time to sample from $\eta_n^{(\nu,s)}$ is in $\mathcal{O}(n^d)$ for some $d > 0$ then the expected time to complete Algorithm 1 is finite whenever $d < b - 1$.*

While one can sample from $\eta_n^{(\nu,s)}$ in principle using either of the algorithms in Section 3.2 and an appropriate sequence of i-cSMC kernels $(P_N^{(n)})_{n\geq 1}$ with corresponding constants $(\beta_n)_{n\geq 1}$, defining these sequences in practice may constitute a research programme on their own. The problem of being able to sample according to (8) in finite expected time is a separate issue which has already been addressed for a few Markov kernels in Flegal and Herbei (2012).

An alternative approach to Algorithm 1 when one can sample according to (8) in finite expected time, but is interested only in unbiased estimation of $\pi(f)$ for some $f : \mathsf{X} \to \mathbb{R}$ (as opposed to perfect simulation) also follows from viewing $\eta_n^{(\nu,s)}$ as a marginal distribution associated with a Feynman–Kac path measure. It is well established that SMC algorithms such as Algorithm 7 produce unbiased estimates of unnormalized expectations $\gamma_n(f)$ (Del Moral, 2004, Section 7.4.2), where γ_n is defined in Section 5.1. In this context, this means that one can unbiasedly estimate $\gamma_n(f) = \mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)\eta_n(f)$ using Algorithm 7. From the decomposition (6), we thus have the identity

$$\pi(f) = \sum_{k \geq 1} \frac{\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq k)}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})} \eta_k(f) = \frac{1}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})} \sum_{k \geq 1} \gamma_k(f).$$

Letting g be a p.m.f. with support $\mathbb{N}$, Algorithm 10 provides an unbiased estimate of $\pi(f)$. In practice, one should choose g so that it has heavier tails than (8). The validity of the method is immediate upon viewing $\mathbb{I}(M = 1)$ as a Bernoulli random variable with success probability

$$\frac{\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq 1)}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})} = \frac{1}{\mathbb{E}_{\nu,s}(\tau_{\nu,s})},$$

and $Z/g(K)$ as an importance sampling estimate of $\sum_{k \geq 1} \gamma_k(f)$. The relative variance of $\mathbb{I}(M = 1)$ is therefore $\mathbb{E}_{\nu,s}(\tau_{\nu,s}) \left[1 - \mathbb{E}_{\nu,s}(\tau_{\nu,s})^{-1}\right]$ and that of $Z/g(K)$ will depend on both g and the number of particles N used to estimate $\gamma_K(f)$ (Del Moral, 2004, Section 9.4.1).

Algorithm 10 Unbiased estimation of $\pi(f)$ using SMC

1. Simulate M from the distribution with p.m.f. (8). If $M \neq 1$, output 0.
2. Sample $K \sim g$ and compute Z, an unbiased estimate of $\gamma_K(f)$ using SMC. Output $Z/g(K)$.

6.5 Parallel implementation of Markov chain Monte Carlo

Algorithms that can produce perfect samples in finite expected time can immediately be implemented in parallel. However, simulation of the split chain $\tilde{\mathbf{X}}_{\nu,s}$ is all that is required for parallel implementation of MCMC, as noted by Mykland et al. (1995). Indeed, if one can detect regeneration events, then one can simulate multiple i.i.d. tours of the split chain in parallel and piece them together afterwards. Brockwell and Kadane (2005) introduced artificial atoms for exactly this purpose, since detection of regeneration is trivial for the split chain $\tilde{\mathbf{X}}_{a,p}$.

Judicious choices of Markov kernel incorporating an artificial atom are therefore highly pertinent in the current computing landscape, where parallel algorithms offer significant computational advantages (see, e.g., Lee et al., 2010). This type of regenerative parallelization is complementary to parallel implementation of a sampling from a complex Markov transition kernel, which typically require the use of a multiple cores on a single computer, as it allows tours of such Markov chains to be simulated on different computers. Notably, and in contrast to the perfect simulation algorithms described in Section 3.2, there is no need for explicit knowledge of bounds on $p(x)$ as one can simply use $s(x) = p(x)$ and simulate the split chain described in Section 3.1.

We now show that Markov chains possessing an identifiable atom can be suitable for parallel implementation in the manner suggested without stringent requirements on their rates of convergence. This is important because in some cases, complex and intractable Markov kernels used in statistical applications are not geometrically ergodic (see, e.g., Andrieu and Roberts, 2009; Andrieu and Vihola, 2012; Lee and Łatuszyński, 2012). Since the lengths of each tour are i.i.d. random variables, we consider here the length of the first tour $\tau_{a,p}$. From Kac's theorem, we have $\mathbb{E}_{a,p}(\tau_{a,p}) = \pi(\{a\})^{-1} < \infty$, and for performance this suggests that having $\pi(\{a\})$ not too close to 0 or 1 will ensure that at least the expected length of a tour is small. In a parallel setting, however, it is important that the expected length of the longest of n independent tours be reasonably short, and this can be guaranteed in some sense when $\mathbb{E}_{a,p}\left(\tau_{a,p}^2\right) < \infty$.

It follows indeed from David and Nagaraja (2003, Section 4.2) that when $\mathbb{E}_{a,p}\left(\tau_{a,p}^2\right) < \infty$,

$$\mathbb{E}_{a,p}\left[\max\left\{\tau_1^{(a,p)}, \ldots, \tau_n^{(a,p)}\right\}\right] \leq \mathbb{E}_{a,p}\left(\tau_{a,p}\right) + (n-1)\left[\frac{\mathbb{E}_{a,p}\left(\tau_{a,p}^2\right) - \mathbb{E}_{a,p}\left(\tau_{a,p}\right)^2}{2n-1}\right]^{1/2},$$

and so one can be assured that the expected length of the longest tour grows at most as $\mathcal{O}(n^{1/2})$. In cases where the tours are geometrically distributed with mean ϵ^{-1}, one can have much slower growth rates since then (see, e.g, Eisenberg, 2008)

$$\frac{1}{\lambda}H_n \leq \mathbb{E}_{a,p}\left[\max\left\{\tau_1^{(a,p)}, \ldots, \tau_n^{(a,p)}\right\}\right] \leq 1 + \frac{1}{\lambda}H_n,$$

where $\lambda = -\log(1-\epsilon)$ and $H_n = \sum_{k=1}^{n} k^{-1}$ is the nth harmonic number. Since H_n grows as $\mathcal{O}(\log n)$, one can expect very few long tours in these situations.

The following proposition, which is essentially a corollary of Bednorz et al. (2008, Theorem 1.1), implies that if $\mathbf{X}$ is a reasonable Markov chain in a specific sense then $\mathbb{E}_{\nu,p}\left(\tau_{\nu,p}^2\right) < \infty$.

Proposition 14. *Let $\mathbf{X}$ be a time-homogeneous, Harris recurrent and aperiodic Markov chain with unique stationary distribution π. If $\mathbf{X}$ possesses a proper atom α, then a $\sqrt{n}$-CLT holds for all bounded functions $g : \mathsf{X} \to \mathbb{R}$ if and only if $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty$.*

Remark. That $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty$ implies a $\sqrt{n}$-CLT holds for all bounded functions $g : \mathsf{X} \to \mathbb{R}$ is provided by Jarner and Roberts (2002, Theorem 4.3).

In particular, Proposition 14 implies that if a $\sqrt{n}$-CLT holds for all bounded functions, then the variance of the return time to α is finite. In a statistical application, where π is a posterior distribution, this suggests that if the chain $\mathbf{X}$ is suitable for estimating any posterior probability then $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right)$ is finite. As an example, if $\mathbf{X}$ is polynomially ergodic and Π satisfies jointly, for some function $V : \mathsf{X} \to [1, \infty)$, constants $c, b \in (0, \infty)$ and $\gamma \in [1/2, 1)$,

$$\int_{\mathsf{X}} V(y)\Pi(x, \mathrm{d}y) \leq V(x) - cV(x)^{\gamma} + b\mathbf{1}_{\alpha}(x), \quad x \in \mathsf{X},$$

then $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty$ (Jarner and Roberts, 2002, Section 4).

6.6 Remarks

The main limitation of the proposed methodology is the requirement that β be chosen to be less than or equal to $\underline{p}$. Nevertheless, the ability in principle to obtain perfect samples from the path distribution associated with a Feynman–Kac model is encouraging, as no previous algorithm was able to do so with guaranteed expected computational time polynomial in the time horizon of the model. We anticipate future progress in at least two areas: deterministic lower bounds on $\underline{p}$ under reasonably weak assumptions and quantitative bounds on the probability that $\beta > \underline{p}$ under suitable assumptions on how Π is defined, how β is chosen and the results of any diagnostics. As an example, a detailed analysis of the suggested i-cSMC kernel construction of Section 5 may provide some assurances in practical applications.

An alternative perfect simulation algorithm has been proposed in Andrieu et al. (2012) to sample from Feynman–Kac path measures. Their algorithm relies on a dominated CFTP procedure (Kendall, 2004) and variants of Algorithm 7 and Algorithm 9 which use a branching mechanism instead of a multinomial resampling mechanism to limit particle interactions. This algorithm has been successfully applied to the simulation of self-avoiding random walks, but computational guarantees have not yet been established. The approaches developed here and in Andrieu et al. (2012) are complementary.

The introduction of an artificial atom may even be beneficial in cases where an atom does exist for Π. Consider the case where $\pi(\{a\}) \leq \zeta$ for any $a \in \mathsf{X}$ in which case it follows that $\inf_{x \in \mathsf{X}} \Pi(x, \{a\}) \leq \zeta$ for any $a \in \mathsf{X}$. If ζ is small then the perfect simulation procedures we have provided are computationally expensive. Furthermore, even if an atom is used only for parallel implementation of MCMC as in Section 6.5, the tours can be prohibitively long when ζ is small since the expected return time to an accessible atom a is $\pi(\{a\})^{-1}$ by Kac's Theorem.

7 Applications

For simplicity, all of our applications are associated with discrete-time, time-homogeneous HMMs with $\mathsf{Z} = \mathbb{R}$. That is, each model is determined by the initial distribution of the latent states μ and for each $p \in \{2, \ldots, n\}$, $M_p = M$. The observations are encoded in the potential functions as $G_p(z_p) := g(z_p, y_p)$, the likelihood of observation $Y_p = y_p$ conditional on $Z_p = z_p$.

In the first two applications, we make use of Corollary 3 in Appendix B. In particular, we will show that for some $A < \infty$,

$$\sup_{z_{p-1} \in S_{p-1}, z_p \in S_p} \frac{G_p(z_p)}{\int_{\mathsf{Z}} G_p(z'_p) M_p(z_{p-1}, \mathrm{d}z'_p)} \cdot \sup_{z_p, z'_p \in S_p, z_{p+1} \in S_{p+1}} \frac{M_{p+1}(z_p, \mathrm{d}z_{p+1})}{M_{p+1}(z'_p, \mathrm{d}z_{p+1})} \leq A, \tag{28}$$

where $S_p := \{z \in \mathsf{Z} : G_p(z) > 0\}$ for each $p \in \{1, \ldots, n\}$, with $S_0 := \mathsf{Z}$. This allows us to take $F \leq A$ in (23). As suggested in Section 5.4, we then take $N = 7(A-1)$ to guarantee that $\epsilon_N \geq 0.75$.

Whenever (28), or the more general condition (30) in Appendix B, can be established in such a way that A is independent of n then it follows that one only needs a number of particles N linear in n in order to control ϵ_N. As the N particles need to be propagated for n time steps in Algorithm 9 the computational cost per simulation from the i-cSMC Markov kernel is $\mathcal{O}(n^2)$. Since ϵ_N has been controlled, the expected number of such simulations required to obtain a perfect sample is bounded by a constant depending on β, and so the overall computational complexity is $\mathcal{O}(n^2)$.

7.1 Particle in an absorbing medium

We let μ be the uniform distribution on $[0, 1]$ and $M(z, \cdot)$ be the Markov kernel associated with a normal distribution with mean z and variance σ^2, and whose density w.r.t. the Lebesgue measure is denoted by $M(z, z')$. For some $S \in \mathcal{B}(\mathsf{X})$, we let $G_p(z_p) = \mathbf{1}_S(z_p)$ for each $p \in \{1, \ldots, n\}$. Under the dynamics associated with μ and M, π is the distribution of the path of a particle that has not left the set S. This is a simple example of a "hard obstacle" model, which has been studied in Del Moral and Doucet (2004). The applicability of Algorithm 7 to obtain accurate estimates of $\pi(f)$ for this model has been discussed in Del Moral and Doucet (2014), where it is also remarked that rejection sampling using the Markov chain defined by μ and M alone has an expected computational cost that is $\gamma_n(1)^{-1}$, which is exponentially increasing in n. Algorithm 7, however, does not produce samples exactly distributed according to π.

As an example, we take $S = [0, 1]$ and $\sigma^2 = 0.25$, from which we can see that (28) is satisfied with

$$A = \frac{1}{M(0, [0, 1])} \frac{M(1, 1)}{M(0, 1)} < 15.5.$$

Since $\gamma_n(1) \leq M(0.5, S)^n = (0.683)^n$ it is clear that for $n = 100$ the expected number of samples for a rejection sampler exceeds 10^{16}. In contrast, $F \leq A$ suggests one can use $7(A-1)n \approx 10^4$ particles to ensure $\epsilon_N > 0.75$. As n grows, the difference becomes more and more significant; for example when $n = 1000$, the expected number of paths sampled in the rejection sampler is around 10^{165} whereas $7(A-1)n \approx 10^5$.

As suggested in Section 5.6, Algorithm 7 was run to estimate the parameters $(\psi_p)_{p \in \{1, \ldots, n\}}$ and b was set to 0.5. A number of simulations indicated that $\check{P}_N((0.5, \ldots, 0.5), \alpha_n)$ was very likely to be greater than 0.47, and we chose $\beta = 0.2$, and $\epsilon = 0.1$ to be conservative.

We ran both algorithms to obtain perfect samples from π with $n = 100$. The number of simulations from $\check{P}_N$ was in perfect agreement with Propositions 3 and 4. In our simulations this average cost of flipping a $(1-p)/(1-\epsilon)$-coin in terms of simulations from $\check{P}_N$ was around 5.5 on average. This indicates an expected number of simulations from $\check{P}_N$ to obtain a perfect sample from $\check{\pi}$ of around 65 and to obtain a perfect sample from π of about 130, but it is worth noting that this figure depends only on β and ϵ and not n. Of course, a less conservative estimate of β would bring this

number lower. The diagnostic procedure of Section 6.1 was run at each state visited by the Markov chain without issue.

Empirically, it seems that a smaller value $N = 10^4$ was sufficient to ensure $\inf_{x \in \check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$ greater than 0.4 and it is true more generally that bounds such as (28) are naturally conservative. We note that if one was to consider a larger set S or a smaller variance σ^2 for the Markov kernels in $\mathcal{M}$ one could obtain smaller values of C by using Lemma 2 with a value of m greater than 1.

7.2 Interval-censored sensor data

We consider a situation in which there are evenly spaced sensors which track an object moving in $\mathbb{R}$. Each sensor $j \in \mathbb{Z}$ is associated with the interval of $[j, j+1)$ of $\mathbb{R}$. The observations are thus integer valued and $Y_p = y_p$ indicates that $Z_p \in [y_p, y_{p+1})$. We let μ be a standard normal distribution and $M(z, \cdot)$, as in Section 7.1, be the Markov kernel associated with a normal distribution with mean z and variance σ^2. In this case, we obtain a bound on A and therefore C that is dependent on the data observed.

It is clear that $S_p := [Y_p, Y_p + 1)$ for each $p \in \{1, \ldots, n\}$.

Assume $Y_p = i$ and $Y_{p+1} = j$. Then

$$\sup_{z_p, z_p' \in S_p, z_{p+1} \in S_{p+1}} \frac{M(z_p, \mathrm{d}z_{p+1})}{M(z_p', \mathrm{d}z_{p+1})} \leq \begin{cases} \frac{M(i+1,j+1)}{M(i,j+1)} & i \leq j, \\ \frac{M(i,j)}{M(i+1,j)} & i > j. \end{cases}$$

Similarly, we have

$$\inf_{z_{p-1} \in S_{p-1}} M_p(G_p)(z_{p-1}) = \begin{cases} M(i, S_p) & i \leq j, \\ M(i+1, S_p) & i > j. \end{cases}$$

and $\bar{G}_p = \bar{G} = 1$. Then, given observations $(Y_1, \ldots, Y_n) = (y_1, \ldots, y_n)$, we bound A by taking the two consecutive observations y_k and y_{k+1} that are furthest apart. We consider the case where $y_k = i$ and $y_{k+1} = j$ and $i < j$, the cases with $i \geq j$ being analogous. Then we set

$$A = \frac{1}{M(i, [j, j+1))} \frac{M(i+1, j)}{M(i, j+1)}.$$

As above, Algorithm 7 was run to estimate the parameters $(\psi_p)_{p \in \{1,\ldots,n\}}$ and b was set to 0.5. A number of simulations indicated that $\inf_{x \in \check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$ was very likely to be greater than 0.5, and we chose $\beta = 0.2$, and $\epsilon = 0.1$ to be conservative. We ran both algorithms to obtain perfect samples from $\check{\pi}$ with $n = 100$ on a simulated data set where $\max_{p \in \{1,\ldots,n-1\}} |y_p - y_{p+1}| = 3$. This led to $A = 38$ with $\sigma^2 = 5$ and therefore $N = 7(A-1)n = 25900$.

The number of simulations from $\check{P}_N$ was again in perfect agreement with Propositions 3 and 4, with the average number of flips of a p-coin to obtain a flip of a $(1-p)/(1-\epsilon)$-coin being just over 5. This indicates an expected number of simulations from $\check{P}_N$ to obtain a perfect sample from π of a little over 120, and again this figure depends only on β and ϵ and not n. Of course, a less conservative estimate of β would bring this number lower. The diagnostic procedure of Section 6.1 was run at each state visited by the Markov chain without issue.

Our methodology is not only suited to indicator observations, but we have used them here for simplicity. In addition, this is a setting where Z is not compact, but each π_p is compactly supported with a support that is data dependent. This type of compact support assumption is very much tied to obtaining bounds on F via (28).

7.3 Linear Gaussian state space model

In practice, (28) can fail to be satisfied for any finite A, even though F is finite whenever the potentials are bounded. One such case is the linear Gaussian model, where μ is Gaussian, $M(z, \cdot)$

is Gaussian with a mean depending linearly on z and a variance independent of z and $g(z, \cdot)$ is Gaussian with a mean depending linearly on z and a variance independent of z. In such cases, π is tractable and perfect samples may be obtained from a Kalman smoother.

We simulated a time series of length $n = 100$ in the case where $M(z, z') = \mathcal{N}(z'; 0.9z, 1)$ and $g(z, y) = \mathcal{N}(y; z, 1)$. After running Algorithm 7 with 10000 particles to estimate the parameters $(\psi_p)_{p\in\{1,\ldots,n\}}$ and setting b to 0.5, it was verified empirically that $\inf_{x\in\check{\mathsf{X}}} \check{P}_N(x, \alpha_n)$ was likely to be greater than 0.5 with $N = 4096$. In addition, comparison of each ψ_p with the quantity $\gamma_p(1)/\gamma_{p-1}(1)$ computed using the Kalman filter showed relative discrepancies of less than 0.02 and $\gamma_n(1)^{-1}\prod_{p=1}^{n}\psi_p = 1.01$ suggested the parameter settings were successful.

Running both perfect simulation algorithms with $\beta = 0.2$ and $\epsilon = 0.1$, along with the diagnostic scheme in Section 6.1, suggested that the samples obtained were perfect. Comparison of the samples with π, obtained from the Kalman smoother, also did not suggest any bias. As before, the computational expense was in perfect agreement with Propositions 3 and 4 and on average less than 6 p-coin flips were needed to obtain a flip of a $(1-p)/(1-\epsilon)$-coin.

7.4 Parallel particle marginal Metropolis–Hastings

Our final application involves not perfect simulation, but parallel implementation of the particle marginal Metropolis–Hastings (PMMH) Markov chain introduced in Andrieu et al. (2010). This Markov chain is well-suited to estimating static parameters of an HMM, and we use the same model and data as in Section 7.3.

Denoting by $\Theta = \mathbb{R}^4$ the parameter space with $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$ governing M and g by $M(z, z') = \mathcal{N}(z'; \theta_1 z, \theta_2)$ and $g(z, y) = \mathcal{N}(y; \theta_3 z, \theta_4)$. The distribution π of interest is the posterior distribution of θ conditional upon the observed data, and can be written as

$$\pi(\theta) := \frac{\gamma_{\theta,n}(1)\varpi(\theta)}{\int \gamma_{\theta',n}(1)\varpi(\theta')\mathrm{d}\theta'},$$

where ϖ is a prior density for θ, and $\gamma_{\theta,n}$ is the measure defined by (18) for a given value of the parameter θ. More specifically, instead of a single μ, $\mathcal{M}$ and $\mathcal{G}$ there are now collections $\{\mu_\theta : \theta \in \Theta\}$, $\{\mathcal{M}_\theta : \theta \in \Theta\}$, and $\{\mathcal{G}_\theta : \theta \in \Theta\}$, and for each θ the triplet $(\mu_\theta, \mathcal{M}_\theta, \mathcal{G}_\theta)$ specifies a Feynman–Kac model. We denote by $G_{\theta,p}$ the time p potential specified by $\mathcal{G}_\theta$. In this context, since the Feynman–Kac models specify different versions of an HMM, the quantity $\gamma_{\theta,n}(1)$ is equal to the marginal likelihood of the observations conditional upon the parameter θ.

While $\gamma_{\theta,n}(1)$ cannot be computed in practice, the PMMH Markov chain is an example of a pseudo-marginal Markov chain (Beaumont, 2003; Andrieu and Roberts, 2009) that in this context can be viewed as evolving on $\Theta \times \mathbb{R}_+$ with transition kernel

$$\Pi(\theta, w; \mathrm{d}\theta', \mathrm{d}w') := \min\left\{1, \frac{\varpi(\theta')q(\theta', \theta)w'}{\varpi(\theta)q(\theta, \theta')w}\right\} U_{\theta'}^N(\mathrm{d}w')q(\theta, \mathrm{d}\theta') + r(\theta, w)\delta_{(\theta,w)}(\mathrm{d}\theta', \mathrm{d}w'),$$

where

$$r(\theta, w) := 1 - \int_{\Theta\times\mathbb{R}_+} \min\left\{1, \frac{\varpi(\theta')q(\theta', \theta)w'}{\varpi(\theta)q(\theta, \theta')w}\right\} q(\theta, \mathrm{d}\theta')U_{\theta'}^N(\mathrm{d}w'),$$

and U_θ^N is the distribution of the random variable

$$\gamma_{\theta,n}^N(1) := \prod_{p=1}^{n} \frac{1}{N}\sum_{j=1}^{N} G_{\theta,p}(\zeta_p^j), \tag{29}$$

when $(\zeta_1^1, \ldots, \zeta_n^N) \sim Q_\theta^N$, i.e., produced by running Algorithm 7 with the Feynman–Kac model specific to θ. The marginal distribution of the θ-coordinate of the invariant distribution associated with this Markov kernel is π. We consider the case where $q(\theta, \cdot)$ is a multivariate normal distribution centred at θ with variance-covariance matrix $0.1^2 Id$, ϖ is a multivariate normal distribution centred at 0 with variance-covariance matrix $0.2^2 Id$, and $N = 2048$ to define U_θ^N.

We modify Π to produce $\check{\Pi}$ according to the first strategy in Section 4.2. We first identify a central parameter θ^*, here taken to be $(0.9, 1, 1, 1)$ for simplicity, and define the artificial atom to be $\alpha := \{a\}$, $a := (a_\theta, a_w)$, where a_θ is a distinguished point and a_w is an estimate of $\gamma_{\theta^*,n}(1)$ obtained by running Algorithm 7 and computing (29) with N large. The prior density for θ is modified to $\check{\varpi}(\theta) = 0.5\mathbb{I}(\theta \in \Theta)\varpi(\theta) + 0.5\mathbb{I}(\theta = a_\theta)$. The "re-entry" proposal μ in (16) is taken to be $\mu(\mathrm{d}\theta, \mathrm{d}w) = q(\theta^*, \mathrm{d}\theta)U_\theta^N(\mathrm{d}w)$, and the mixture weights in (16) are taken to be 0.5. Simulation of the split chain with $\nu = \delta_a$ and $s = p$ associated with the transition kernel $\check{\Pi}$ is then straightforward.

We ran the modified Markov chain $\check{\mathbf{X}}$ for 100000 iterations, obtaining 71713 tours. The proportion of time spent at the artificial atom was therefore 0.72 and the average length of a tour was 1.4. The empirical variance of the tour lengths was 7.7 and the length of the longest tour was 150. There were only 5 tours whose length exceeded 100. It is possible that the value of a_w was slightly too large, as almost three quarters of the samples of $\check{\mathbf{X}}$ are discarded when $\check{\mathbf{Y}}$ is constructed. However, this factor of 4 loss is offset by the potential for a practically arbitrary gain in efficiency by distributing the simulation of multiple tours across a number of different computers.

Acknowledgements

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme *Advanced Monte Carlo Methods for Complex Inference Problems*, where work on this paper was undertaken. Arnaud Doucet's research is supported by an Engineering and Physical Sciences Research Council Established Career Research Fellowship.

A Proofs

A.1 Proofs for Section 3

Proof of Proposition 1. Since $p > \epsilon > 0$, we have the Maclaurin series expansion $1/p = \sum_{k=0}^{\infty}(1-p)^k$ and so

$$\frac{\epsilon}{p} = \epsilon\sum_{k=1}^{\infty}(1-p)^{k-1} = \sum_{k=1}^{\infty}\epsilon(1-\epsilon)^{k-1}\left(\frac{1-p}{1-\epsilon}\right)^{k-1},$$

which is an expectation w.r.t. a Geometric(ϵ) random variable, and the result follows. □

Proof of Proposition 2. We define the stopping time $\tau := \inf\{n \geq 1 : Y_n + Z_n \geq 1\}$ and it is clear that

$$\begin{aligned}\mathbb{P}(Y_\tau = 1) &= \mathbb{P}(Y_n = 1 \mid Y_n + Z_n \geq 1)\\ &= \frac{\epsilon}{\epsilon + (p-\epsilon)/(1-\epsilon) - \epsilon(p-\epsilon)/(1-\epsilon)} = \frac{\epsilon}{p}.\end{aligned}$$

We have $\mathbb{E}(\tau) = [1-(1-\epsilon)(1-p)/(1-\epsilon)]^{-1} = p^{-1}$ since τ is the minimum of two geometric random variables and hence also geometrically distributed. Now when $Y_\tau = 1$ it is not necessary to flip Z_τ, and so the expected number of $(p-\epsilon)/(1-\epsilon)$-coin flips is

$$\frac{\epsilon}{p}\left(\frac{1}{p} - 1\right) + \left(1 - \frac{\epsilon}{p}\right)\frac{1}{p} = \frac{1-\epsilon}{p}.$$

□

Proof of Corollary 1. We note that simulating X_n and $\rho_{n-1}^{(a,\epsilon)}$ requires drawing $X_n \sim \Pi(X_{n-1}, \cdot)$ and then, if $X_n = a$, flipping an $\epsilon/p(X_{n-1})$-coin. From $p(X_{n-1}) = \Pi(X_{n-1}, \{a\})$ and Proposition 2 we have that the expected number of $(1-p(X_{n-1}))/(1-\epsilon)$-coin flips required to simulate $(\rho_{n-1}^{(a,\epsilon)}, X_n)$ is $p(X_{n-1})(1-\epsilon)/p(X_{n-1}) = 1-\epsilon$. □

Proof of Proposition 3. For Algorithm 4, the number of samples from Π is ϵ^{-1} since we simulate $X_2, \ldots, X_{\tau_{a,\epsilon}+1}$ each using Π and the expected length of a tour is ϵ^{-1}. The expected number of $(1-p)/(1-\epsilon)$-coin flips required is $(1-\epsilon)/\epsilon$ from Proposition 2, the fact that we simulate $\rho_1^{(a,\epsilon)}, \ldots, \rho_{\tau_{a,\epsilon}}^{(a,\epsilon)}$ and the expected length of each tour being ϵ^{-1}.

For Algorithm 5, the number of $(1-p)/(1-\epsilon)$-coin flips required is $\epsilon^{-1}-1$ since $N \sim \text{Geometric}(\epsilon)$. For each of the samples $X_2, \ldots, X_N$ from $R_{a,\epsilon}(x, \cdot)$ the probability of selecting the mixture component $R_{a,p}$ in (12) is $(1-p(x))/(1-\epsilon)$ and the expected number of samples from $\Pi(x, \cdot)$ to obtain a sample from $R_{a,p}(x, \cdot)$ by rejection is $1/(1-p(x))$. Hence the expected number of samples from Π for each of the $N-1$ steps is ϵ^{-1}. □

The following Lemma is used to prove Proposition 4.

Lemma 1. *Let $B \geq 0.5$, $b \in [B, 1]$, $p \in [0, b]$ and $C = 2/(1+b)$, with $\varepsilon = 1 - Cb$. Then the expected number of p-coin flips required to flip a Cp-coin using the algorithm of Huber (2014) with parameters (C, ε) is bounded above by 11.*

Proof. From inspection of the default algorithm settings, one can check that since $B \geq 0.217$ and $b \geq B$, one indeed has $\varepsilon = \min\{0.644, 1 - Cb\} = 1 - Cb$. From Huber (2014, Theorem 3.6), we have that the number of p-coin flips T satisfies

$$\begin{aligned}
\mathbb{E}[T] &\leq \frac{k(C-1)+C}{1-(Cp)^k} - \frac{C-1}{1-Cp} + \frac{\frac{r}{1-r}[\gamma k(\frac{C}{1-\varepsilon}-1) + (1-\gamma)^2\frac{C}{1-\varepsilon}]}{1-(Cp)^k} \\
&\leq \frac{1}{1-(Cp)^k}\left\{k(C-1) + C + \frac{r}{1-r}\left[\gamma k(\frac{C}{1-\varepsilon}-1) + (1-\gamma)^2\frac{C}{1-\varepsilon}\right]\right\}
\end{aligned}$$

where $k := 2.3/(\gamma\varepsilon)$ and $r := \exp(-2.3)/(1-\gamma)^2$ and $\gamma = 1/2$. We first bound

$$(Cp)^k \leq (Cb)^k = \left(\frac{2b}{1+b}\right)^k \nearrow \exp(-4.6)$$

as $b \to 1$ so $\frac{1}{1-(Cp)^k} \leq \frac{1}{1-\exp(-4.6)} < 1.0102$. Noting that $\varepsilon = C - 1$ we have $k(C-1) + C = 4.6 + 2/(1+b)$. Finally, since $C/(1-\varepsilon) = b^{-1}$,

$$\gamma k\left(\frac{C}{1-\varepsilon} - 1\right) = \frac{2.3}{C-1}\left(\frac{1-b}{b}\right) = 2.3\frac{1+b}{b},$$

and $(1-\gamma)^2 C/(1-\varepsilon) = 1/(4b)$. All together, we obtain

$$\mathbb{E}[T] \leq 1.0102\left\{4.6 + \frac{2}{1+b} + \frac{r}{1-r}\left[2.3\frac{1+b}{b} + \frac{1}{4b}\right]\right\},$$

which is monotonically decreasing in b. Since $r/(1-r) \leq 0.6696$ we have

$$\mathbb{E}[T] \leq 1.0102\left\{4.6 + \frac{2}{1+B} + 0.6696\left[2.3\frac{1+B}{B} + \frac{1}{4B}\right]\right\},$$

and for $B = 0.5$ we obtain $\mathbb{E}[T] \leq 11$. □

Remark. A bound can be obtained for any $B \geq 0.217$ using the final displayed equation. One could potentially derive better choices of γ and k based on the developments in the proof.

Proof of Proposition 4. This follows from Lemma 1 by noting that

$$C = \frac{1}{1-\epsilon} = \frac{1}{1-\beta/2} = \frac{2}{2-\beta} = \frac{2}{1+b},$$

with $b = 1 - \beta \geq 0.5$, and $q = 1 - p \leq 1 - \beta = b$. □

A.2 Proof for Section 4

Proof of Proposition 5. Since $\mathbf{X}$ is uniformly ergodic, Π satisfies for some $m \in \mathbb{N}$, $\epsilon > 0$ and probability measure ν,

$$\Pi^m(x, A) \geq \epsilon\nu(A), \quad A \in \mathcal{B}(\mathsf{X}).$$

We consider $\check{\Pi}^{m+1}$ and an arbitrary $A \in \mathcal{B}(\mathsf{X})$. It follows that if $x \in \mathsf{X}$ then from (17),

$$\check{\Pi}^{m+1}(x, A) \geq \int_{\mathsf{X}} \check{\Pi}(x, \mathrm{d}y)\check{\Pi}^m(y, A) \geq w^{m+1}\epsilon\nu(A),$$

since $\check{\Pi}(x, \mathsf{X}) \geq w\Pi(x, \mathsf{X}) = w$, and

$$\check{\Pi}^{m+1}(a, A) \geq \int_{\mathsf{X}} \check{\Pi}(a, \mathrm{d}y)\check{\Pi}^m(y, A) \geq \check{\Pi}(a, \mathsf{X})w^m\epsilon\nu(A).$$

Hence $\inf_{x\in\check{\mathsf{X}}} \check{\Pi}^{m+1}(x, A) \geq \min\left\{w, \check{\Pi}(a, \mathsf{X})\right\} w^m\epsilon\nu(A)$ and, as $\check{\Pi}(a, \mathsf{X}) > 0$ by assumption, it follows that $\check{\mathbf{X}}$ is uniformly ergodic.

To show that the converse does not hold, consider the case where $\check{\Pi} = w\Pi_1 + (1-w)\Pi_2$ is exactly as in the mixture strategy discussed in Section 4.2. Then if $C := \inf_{x\in\mathsf{X}} \mu(\mathrm{d}x)/\pi(\mathrm{d}x) > 0$, we have for $x \in \mathsf{X}$

$$\check{\Pi}(x, \{a\}) \geq (1-w)\min\left\{1, \check{\pi}(\{a\})\frac{\mu(\mathrm{d}x)}{\check{\pi}(\mathrm{d}x)}\right\} \geq (1-w)C\check{\pi}(\{a\})$$

and $\check{\Pi}(a, \{a\}) \geq w$. Hence, $\check{\Pi}(x, \{a\}) \geq \min\left\{w, (1-w)C\check{\pi}(\{a\})\right\} > 0$ and so $\check{\mathbf{X}}$ is uniformly ergodic, irrespective of the properties of Π. □

A.3 Proofs for Section 5

Proof of Proposition 6. Since the potentials are strictly positive, the measures $Q^N(\mathrm{d}x, \mathrm{d}v)$ and $\pi(\mathrm{d}x)\bar{Q}_x(\mathrm{d}v)$ are equivalent. Therefore,

$$\begin{aligned}
Q^N(\mathrm{d}x) &= \sum_{\mathbf{k}\in[N]^n} \int_{\mathsf{V}_N} Q^N(\mathbf{k}, \mathrm{d}x, \mathrm{d}v) \\
&= \sum_{\mathbf{k}\in[N]^n} \int_{\mathsf{V}_N} \frac{Q^N(\mathbf{k}, \mathrm{d}x, \mathrm{d}v)}{\pi(\mathrm{d}x)\bar{Q}_x^N(\mathbf{k}, \mathrm{d}v)}\pi(\mathrm{d}x)\bar{Q}_x^N(\mathbf{k}, \mathrm{d}v) \\
&= \pi(\mathrm{d}x) \int_{\mathsf{V}_N} \frac{1}{\phi^N(v)}\bar{Q}_x^N(\mathrm{d}v) \\
&\geq \pi(\mathrm{d}x)\left\{\int_{\mathsf{V}_N} \phi^N(v)\bar{Q}_x^N(\mathrm{d}v)\right\}^{-1},
\end{aligned}$$

by Jensen's inequality, and we conclude. □

Proof of Proposition 7. We provide here a simple proof leveraging a result of (Andrieu et al., 2013). Under the assumptions given, condition (A1) in Andrieu et al. (2013) always holds when n is finite and their constant α is equal to F (see also Appendix B). Hence we can apply Andrieu et al. (2013, Proposition 14) to obtain that for any $N \geq 2$

$$\sup_{x,y\in\mathsf{X}} \bar{\mathsf{E}}_{x,y}^N\left[\phi^N(V)\right] \leq \left(1 + \frac{2(F-1)}{N}\right)^n,$$

where $\bar{\mathsf{E}}_{x,y}^N$ denotes expectation w.r.t. the law of a doubly conditional SMC algorithm with two fixed paths $x, y \in \mathsf{X}$, a construction defined in Andrieu et al. (2013). Since

$$\sup_{x\in\mathsf{X}} \bar{\mathsf{E}}_x^N\left[\phi^N(V)\right] \leq \sup_{x,y\in\mathsf{X}} \bar{\mathsf{E}}_{x,y}^N\left[\phi^N(V)\right],$$

we conclude. □

Proof of Proposition 8. Since $\psi_1, \ldots, \psi_n$ are constants, the potential functions $\left\{\check{G}_p : p \in \{1, \ldots, n\}\right\}$ are $\check{\pi}$-essentially bounded above if and only if the potential functions $\{G_p : p \in \{1, \ldots, n\}\}$ are π-essentially bounded above. The result then follows from Andrieu et al. (2013, Theorem 1). □

Proof of Proposition 9. This follows immediately from Lemma 3 in Appendix B □

A.4 Proofs for Section 6

Proof of Proposition 10. We define $(\xi_i)_{i\geq 1}$ by $\xi_i := B_i - \beta$ and $S_n := \sum_{i=1}^n \xi_i$ with $S_0 := 0$. Then $\tau = \inf\{n \geq 1 : S_n > 0\}$. For the first two parts, we can think of $(S_n)_{n\geq 0}$ as a Markov chain on $\mathbb{R}$ (Lamperti, 1960, see also Meyn and Tweedie 2009, Chapter 9) with bounded increments. If $p < \beta$, $(S_n)_{n\geq 0}$ is transient with negative drift and so its probability of return to $\mathbb{R}_+$ is strictly inferior to 1. If $p = \beta$, $(S_n)_{n\geq 0}$ is null recurrent and so its probability of return to $\mathbb{R}_+$ is 1, and its expected return time to $\mathbb{R}_+$ is infinite. Finally, when $p > \beta$, we can apply Wald's equation $\mathbb{E}(S_\tau) = \mathbb{E}(\tau)\mathbb{E}(\xi_1)$, and since $S_\tau \leq 1 - \beta$ by construction and $\mathbb{E}(\xi_1) = p - \beta$, we have $\mathbb{E}(\tau) \leq (1 - \beta)/(p - \beta)$. □

Proof of Corollary 2. For the first part, let $A := \{x \in \mathsf{X} : p(x) < \beta\}$ and by assumption $\pi(A) > 0$. Now let $n = \min\{k : \nu R_{a,\epsilon}^{k-1}(A) > 0\}$, which from the definitions of ν and $R_{a,\epsilon}$ must be finite. Moreover, for any $k \in \{1, \ldots, n\}$, $\nu R_{a,\epsilon}^{k-1}$ is the distribution of X_k using either algorithm, conditional upon the algorithm having not stopped before time k. Letting τ_x denote the stopping time of the diagnostic when $p = p(x)$, it follows that the probability that the algorithm never terminates is greater than or equal to $(1 - \epsilon)^{n-1} \int_A \nu R_{a,\epsilon}^{n-1}(\mathrm{d}x)\mathbb{P}(\tau_x = \infty) > 0$.

For the second part, let $A := \{x \in \mathsf{X} : p(x) = \beta\}$ and $n = \min\{k : \nu R_{a,\epsilon}^{k-1}(A) > 0\}$ as above. Then it follows that the expected number of flips of a p-coin is greater than $(1-\epsilon)^{n-1} \int_A \nu R_{a,\epsilon}^{n-1}(\mathrm{d}x)\mathbb{E}(\tau_x) = \infty$. The third part is immediate. □

Proof of Proposition 11. As in the proof of Proposition 10 we define $(\xi_i)_{i\geq 1}$ by $\xi_i := B_i - \beta$ but now define $S_n := \beta^{-1} \sum_{i=1}^n \xi_i = \sum_{i=1}^n (mB_i - 1)$. It is clear that $(S_n)_{n\geq 1}$ is a Markov chain on the integers that is skip-free to the left and so we can apply Brown et al. (2010, Corollary 1) to obtain

$$\mathbb{P}(\tau < \infty) = 1 - \frac{1 - mp}{1 - p} = \frac{p(m-1)}{1-p}.$$

□

Proof of Proposition 12. Both kernels are uniformly ergodic with $\|\Pi^n(x, \cdot) - \pi\|_{\mathrm{TV}} \leq (1 - \underline{\epsilon})^n$ and $\left\|\tilde{\Pi}^n(x, \cdot) - \tilde{\pi}\right\|_{\mathrm{TV}} \leq (1 - \epsilon)^n$. We have, following Mitrophanov (2005, Theorem 3.2),

$$\left\|\tilde{\Pi}^n(x, \cdot) - \Pi^n(x, \cdot)\right\|_{\mathrm{TV}} \leq \frac{1}{\epsilon} \sup_{x\in\mathsf{X}} \left\|\tilde{\Pi}(x, \cdot) - \Pi(x, \cdot)\right\|_{\mathrm{TV}}.$$

It remains to bound $\left\|\tilde{\Pi}(x, \cdot) - \Pi(x, \cdot)\right\|_{\mathrm{TV}} = \sup_{A\in\mathcal{B}(\mathsf{X})} \left|\Pi(x, A) - \tilde{\Pi}(x, A)\right|$. Letting $a \wedge b$ denote $\min\{a, b\}$ and $A \in \mathcal{B}(\mathsf{X})$, we have

$$\begin{aligned}
\left|\Pi(x, A) - \tilde{\Pi}(x, A)\right| &= \left|\Pi(x, A) - \epsilon\delta_a(A) - (1 - \epsilon)\frac{\Pi(x, A) - \{p(x) \wedge \epsilon\}\,\delta_a(A)}{1 - p(x) \wedge \epsilon}\right| \\
&= \left|\Pi(x, A)\left\{1 - \frac{1-\epsilon}{1 - p(x) \wedge \epsilon}\right\} + \delta_a(A)\left\{\frac{1-\epsilon}{1 - p(x) \wedge \epsilon}\{p(x) \wedge \epsilon\} - \epsilon\right\}\right| \\
&= \left|\Pi(x, A)\left\{\frac{\epsilon - p(x) \wedge \epsilon}{1 - p(x) \wedge \epsilon}\right\} + \delta_a(A)\left\{\frac{p(x) \wedge \epsilon - \epsilon}{1 - p(x) \wedge \epsilon}\right\}\right| \\
&= |\Pi(x, A) - \delta_a(A)|\,\frac{\epsilon - p(x) \wedge \epsilon}{1 - p(x) \wedge \epsilon} \\
&\leq |\Pi(x, A) - \delta_a(A)|\,\frac{\epsilon - \underline{\epsilon}}{1 - \underline{\epsilon}},
\end{aligned}$$

where the last inequality follows because $(a-b)/(1-b)$ is monotonically decreasing as b increases and $\underline{p} \geq \underline{\epsilon}$. Because $\Pi(x,\{a\}) \geq \underline{\epsilon}$ we have $|\Pi(x,A) - \delta_a(A)| \leq 1-\underline{\epsilon}$ and therefore $\left\|\tilde{\Pi}(x,\cdot) - \Pi(x,\cdot)\right\|_{\mathrm{TV}} \leq \epsilon - \underline{\epsilon}$. Finally,

$$\left\|\tilde{\Pi}^n(x,\cdot) - \Pi^n(x,\cdot)\right\|_{\mathrm{TV}} \leq \frac{1}{\epsilon}\sup_{x\in\mathsf{X}}\left\|\tilde{\Pi}(x,\cdot) - \Pi(x,\cdot)\right\| \leq 1 - \underline{\epsilon}/\epsilon$$

and we conclude. □

Proof of Proposition 13. The proof essentially follows the same argument as the corresponding negative results of Asmussen et al. (1992) and Blanchet and Meng (2007). Let $T(n)$ denote the computational time required to simulate from $\eta_n^{(\nu,s)}$. Then the assumptions provide that there exists $C < \infty$ such that the expected computational time of this simulation is

$$\sum_{n\geq 1} T(n)\mathbb{P}_{\nu,s}(\tau_{\nu,s} \geq n)/\mathbb{E}_{\nu,s}(\tau_{\nu,s}) \leq \sum_{n=1}^{\infty} Cn^d n^{-b},$$

and the r.h.s. is finite whenever $d < b-1$. □

Proof of Proposition 14. From Bednorz et al. (2008, Theorem 1.1) we have that a $\sqrt{n}$-CLT holds for a function g if and only if $\mathbb{E}_{\alpha,p}\left[\left(\sum_{k=1}^{\tau_{\alpha,p}} \bar{g}(X_k)\right)^2\right] < \infty$, where $\bar{g} := g - \pi(g)$. We define $\|g\| := \sup_{x\in\mathsf{X}} |g(x)|$.

($\Longleftarrow$)If $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty$ and $\|g\| < \infty$, we have $\|\bar{g}\| \leq 2\|g\|$ and so

$$\mathbb{E}_{\alpha,p}\left[\left(\sum_{k=1}^{\tau_{\alpha,p}} \bar{g}(X_k)\right)^2\right] \leq \mathbb{E}_{\alpha,p}\left[\left(\sum_{k=1}^{\tau_{\alpha,p}} 2\|g\|\right)^2\right] = 4\|g\|^2 \mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty.$$

($\Longrightarrow$) If $\mathbb{E}_{\alpha,p}\left[\left(\sum_{k=1}^{\tau_{\alpha,p}} \bar{g}(X_k)\right)^2\right] < \infty$ for all g such that $\|g\| < \infty$ then clearly this holds for $g(x) = \mathbb{I}\{x \in \alpha\} - \pi(\alpha) = \bar{g}(x)$. Then

$$\begin{aligned}\mathbb{E}_{\alpha,p}\left[\left(\sum_{k=1}^{\tau_{\alpha,p}} \bar{g}(X_k)\right)^2\right] &= \mathbb{E}_{\alpha,p}\left[\left(1 - \tau_{\alpha,p}\pi(\alpha)\right)^2\right] \\ &= \pi(\alpha)^2 \mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) - 1\end{aligned}$$

as $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}\right) = 1/\pi(\alpha)$ from Kac's theorem so it follows that $\mathbb{E}_{\alpha,p}\left(\tau_{\alpha,p}^2\right) < \infty$. □

B Lemmas for F and $\check{F}$

We adopt here the notation of Del Moral (2004). We define for each $p \in \{2,\ldots,n\}$, the non-negative kernel $Q_p(x_{p-1}, \mathrm{d}x_p) := G_{p-1}(x_{p-1})M_p(x_{p-1}, \mathrm{d}x_p)$. We can then define for any $p \in \{2,\ldots,n\}$ and $k \in \{1,\ldots,n-p+1\}$

$$Q_{p,p+k}(A)(z_p) := \int_{\mathsf{Z}^k} \mathbb{I}(z_{p+k} \in A) \prod_{q=p+1}^{p+k} Q_q(z_{q-1}, \mathrm{d}z_q), \quad A \in \mathcal{B}(\mathsf{Z}).$$

This allows us to express, e.g., $\gamma_n(1) = \mu Q_{2,n+1}(1)$. We define for each $p \in \{2,\ldots,n\}$

$$\eta_p(A) := \frac{\mu Q_{2,p}(A)}{\mu Q_{2,p}(1)}, \quad A \in \mathcal{B}(\mathsf{Z}),$$

and the relationship between η_p and π_{p-1} is that $\eta_p(A) = \pi_{p-1}(f)$ where $f(z_1, \ldots, z_{p-1}) = M_p(z_{p-1}, A)$. It follows that $\eta_p(G_p) = \gamma_p(1)/\gamma_{p-1}(1)$ and that $\eta_p Q_{p,p+k}(1) = \gamma_{p+k-1}(1)/\gamma_{p-1}(1)$.

For a Markov kernel P evolving on E, we write $P(f)(x) := \int_{\mathsf{Z}} f(z)P(x, \mathrm{d}z)$. Assumption (A1) of Andrieu et al. (2013) is the existence of a constant $A < \infty$ such that

$$\sup_{z \in \mathsf{Z}} \frac{Q_{p,p+k}(1)(z)}{\eta_p Q_{p,p+k}(1)} \leq A, \qquad p \in \{2, \ldots, n\}, k \in \{1, \ldots, n-p+1\}, \tag{30}$$

and this is in fact exactly F defined in Section 5.4.

Let $S_p := \{z \in \mathsf{Z} : G_p(z) > 0\}$ and $\bar{G}_p := \sup_{z \in \mathsf{Z}} G_p(z)$ for each $p \in \{1, \ldots, n\}$, with $S_0 := \mathsf{Z}$. Let $M_{n+1}(z, A) := \chi(A)$ independent of z and

$$M_{p+1,p+m}(z_p, A) := \int_{\mathsf{Z}^m} \mathbb{I}(z_{p+m} \in A) \prod_{q=p+1}^{p+m} M_q(z_{q-1}, \mathrm{d}z_q), \quad A \in \mathcal{B}(\mathsf{Z}).$$

Lemma 2. *Let $A < \infty$, and assume that for all $p \in \{1, \ldots, n\}$, there exists an $m \in \{1, \ldots, n-p+1\}$ such that*

$$\frac{\prod_{q=p}^{p+m} \bar{G}_q}{\prod_{q=p}^{p+m} \eta_q(G_q)} \cdot \sup_{z_p \in S_p, z_{p+m} \in S_{p+m}} \frac{M_{p+1,p+m}(z_p, \mathrm{d}z_{p+m})}{\eta_{p+m}(\mathrm{d}z_{p+m})} \leq A,$$

then (30) holds with this choice of A.

Proof. If $k \leq m$, then the result holds trivially. For any $k \in \{m+1, \ldots, n-p+1\}$,

$$\begin{aligned}
Q_{p,p+k}(1)(z_p) &= \int_{\mathsf{Z}^k} \prod_{q=p+1}^{p+k} Q_q(z_{q-1}, \mathrm{d}z_q) \\
&= \int_{\mathsf{Z}} Q_{p,p+m}(z_p, \mathrm{d}z_{p+m}) Q_{p+m,p+k}(1)(z_{p+m}) \\
&= \int_{\mathsf{Z}} \frac{Q_{p,p+m}(z_p, \mathrm{d}z_{p+m})}{\left[\prod_{q=p}^{p+m-1} \eta_q(G_q)\right] \eta_{p+m}(\mathrm{d}z_{p+m})} \left[\prod_{q=p}^{p+m-1} \eta_q(G_q)\right] \eta_{p+m}(\mathrm{d}z_{p+m}) Q_{p+m,p+k}(1)(z_{p+m}) \\
&\leq \left[\frac{\prod_{q=p}^{p+m-1} \bar{G}_q}{\prod_{q=p}^{p+m-1} \eta_q(G_q)} \sup_{z_p \in S_p, z_{p+m} \in S_{p+m}} \frac{M_{p+1,p+m}(z_p, \mathrm{d}z_{p+m})}{\eta_{p+m}(\mathrm{d}z_{p+m})}\right] \eta_p Q_{p,p+k}(1).
\end{aligned}$$

□

Corollary 3. *If*

$$\sup_{z_{p-1} \in S_{p-1}, z_p \in S_p} \frac{G_p(z_p)}{M_p(G_p)(z_{p-1})} \cdot \sup_{z_p, z_p' \in S_p, z_{p+1} \in S_{p+1}} \frac{M_{p+1}(z_p, \mathrm{d}z_{p+1})}{M_{p+1}(z_p', \mathrm{d}z_{p+1})} \leq A,$$

then (30) holds with this choice of A.

Proof. The result follows from taking $m = 1$ in Lemma 2 and noting that $\eta_{p+1}(\mathrm{d}z_{p+1}) \geq \inf_{z_p \in S_p} M_{p+1}(z_p, \mathrm{d}z_{p+1})$ and $\eta_p(G_p) \geq \inf_{z_{p-1} \in S_{p-1}} M_p(G_p)(z_{p-1})$. □

When considering the i-cSMC kernel associated with the extended Feynman–Kac model of Section 5.5, it is possible that the choice of A satisfying (30) for the original model does not simultaneously satisfy the associated condition for the extended model. Indeed, one now requires a constant $\check{A} < \infty$ such that

$$\sup_{z \in \mathsf{Z}} \frac{\check{Q}_{p,p+k}(1)(z)}{\check{\eta}_p \check{Q}_{p,p+k}(1)} \leq \check{A}, \qquad p \in \{2, \ldots, n\}, k \in \{1, \ldots, n-p+1\}, \tag{31}$$

where it can be seen that $\check{A}$ is equal to $\check{F}$ defined in Section 5.5.

Lemma 3. *Let $E \geq 1$ be such that for each $p \in \{1, \ldots, n\}$, $E^{-1} \leq \psi_p / \eta_p(G_p) \leq E$. Then (31) holds with $\check{A} = AE^n$, where A satisfies (30).*

Proof. We can write

$$\begin{aligned}
\check{\eta}_p \check{Q}_{p,p+k}(1) &= \check{\eta}_p(\{a\}) \prod_{q=p}^{p+k-1} \psi_q + \check{\eta}_p(\mathsf{Z}) \eta_p Q_{p,p+k}(1) \\
&= \eta_p Q_{p,p+k}(1) \left(\check{\eta}_p(\{a\}) \frac{\prod_{q=p}^{p+k-1} \psi_q}{\eta_p Q_{p,p+k}(1)} + \check{\eta}_p(\mathsf{Z}) \right).
\end{aligned}$$

Now we have for $z = a$

$$\frac{\check{Q}_{p,p+k}(1)(a)}{\check{\eta}_p \check{Q}_{p,p+k}(1)} = \frac{\prod_{q=p}^{p+k-1} \psi_q}{\check{\eta}_p(\{a\}) \prod_{q=p}^{p+k-1} \psi_q + \check{\eta}_p(\mathsf{Z}) \eta_p Q_{p,p+k}(1)} \leq \max \left\{ 1, \frac{\prod_{q=p}^{p+k-1} \psi_q}{\eta_p Q_{p,p+k}(1)} \right\},$$

whereas for $z \in \mathsf{Z}$

$$\begin{aligned}
\frac{\check{Q}_{p,p+k}(1)(z)}{\check{\eta}_p \check{Q}_{p,p+k}(1)} &= \frac{Q_{p,p+k}(1)(z)}{\eta_p Q_{p,p+k}(1)} \left(\check{\eta}_p(\{a\}) \frac{\prod_{q=p}^{p+k-1} \psi_q}{\eta_p Q_{p,p+k}(1)} + \check{\eta}_p(\mathsf{Z}) \right)^{-1} \\
&\leq A \max \left\{ 1, \frac{\eta_p Q_{p,p+k}(1)}{\prod_{q=p}^{p+k-1} \psi_q} \right\}.
\end{aligned}$$

Since $\max \left\{ 1, \frac{\eta_p Q_{p,p+k}(1)}{\prod_{q=p}^{p+k-1} \psi_q} \right\} \leq E^n$ and $\max \left\{ 1, \frac{\prod_{q=p}^{p+k-1} \psi_q}{\eta_p Q_{p,p+k}(1)} \right\} \leq E^n$ the result is immediate. □

C An alternative solution to Problem 2 via a generic solution to the sign problem

We have $(1-p)/(1-\epsilon) = 1 - (p-\epsilon)/(1-\epsilon)$ so simulation of a $(1-p)/(1-\epsilon)$-coin is possible whenever one can simulate a $(p-\epsilon)/(1-\epsilon)$-coin. This in turn is feasible when one can simulate a random variable V satisfying $\mathsf{E}(V) = p$ and $\mathsf{P}(\epsilon \leq V \leq 1) = 1$, since the random variable $\mathbb{I}(U < \frac{V-\epsilon}{1-\epsilon})$ is a $(p-\epsilon)/(1-\epsilon)$-coin flip, where U is an independent uniform random variable on $[0,1]$. We now show that unbiased constrained estimation is possible in a more general sense, and a solution to Problem 2 follows naturally whenever $\epsilon \leq 1/2$.

The first step is a generic solution to what is often referred to as a "sign problem". In particular, given a probability measure $\mu : \mathcal{B}(\mathsf{X}) \to [0,1]$ whose distribution one can obtain any number of i.i.d. samples from, and a function $\varphi : \mathsf{X} \to \mathbb{R}$ such that $\mu(\varphi) > 0$, one wishes to simulate a random variable with expectation $\mu(\varphi)$ that is almost surely non-negative. The non-existence of a general procedure for solving this problem without any assumptions on μ and φ has been shown in Jacob and Thiery (2013) following other non-existence results such as Keane and O'Brien (1994) and Henderson and Glynn (2003). We provide a positive result in the case where φ is a bounded function and $\mu(\varphi) \geq \delta > 0$ for a known constant δ. Our scheme relies on the probability measure

$$\mu_{|\varphi|}(\mathrm{d}x) := \frac{\mu(\mathrm{d}x)\,|\varphi(x)|}{\mu(|\varphi|)},$$

which can be sampled from by rejection whenever $\|\varphi\| := \sup_{x \in \mathsf{X}} |\varphi(x)| < \infty$. A general algorithm for solving the sign problem is provided in Algorithm 11, whose validity is established in Proposition 15.

Proposition 15. *Let φ satisfy $\|\varphi\| < \infty$ and $\mu(\varphi) \geq \delta > 0$. Then the output of Algorithm 11 is a random variable W such that $\mathsf{E}(W) = \mu(\varphi)$ and $\mathsf{P}(0 \leq W \leq \|\varphi\|) = 1$.*

Algorithm 11 Simulate an unbiased and almost surely positive estimate of $\mu(\varphi) \geq \delta > 0$

1. Sample $\xi \sim \mu$.
2. Flip a $2q$-coin Y, where $2q \leq 1 - \delta / \|\varphi\|$, using a Bernoulli factory using flips of a q-coin, each such flip being simulated by sampling $\zeta \sim \mu_{|\varphi|}$ and outputting $\mathbb{I}(\text{sign}(\varphi(\zeta)) = -1)$.
3. Output $|\varphi(\xi)|\,(1 - Y)$.

Proof. We denote the output of Algorithm 11 as $W = |\varphi(\xi)|\,(1 - Y)$. Since $\mu(\varphi) > 0$ implies that $\mu(|\varphi|) > 0$, we can express $\mu(\varphi)$ as

$$\mu(\varphi) = \mu(|\varphi|)\frac{\mu(\varphi)}{\mu(|\varphi|)} = \mu(|\varphi|)\int_{\mathsf{X}} \text{sign}(\varphi(x))\mu_{|\varphi|}(\mathrm{d}x).$$

The random variable W is the product of independent unbiased estimates of $\mu(|\varphi|)$ and $\int_{\mathsf{X}} \text{sign}(\varphi(x))\mu_{|\varphi|}(\mathrm{d}x)$ respectively. That $\mathsf{E}\,|\varphi(\xi)| = \mu(|\varphi|)$ is immediate, and it is clear that $\mathsf{P}(0 \leq |\varphi(\xi)| \leq \|\varphi\|) = 1$. We now show that $\mathsf{E}(1 - Y) = \int_{\mathsf{X}} \text{sign}(\varphi(x))\mu_{|\varphi|}(\mathrm{d}x)$. By construction,

$$\int_{\mathsf{X}} \text{sign}(\varphi(x))\mu_{|\varphi|}(\mathrm{d}x) = \mathsf{E}\left[\text{sign}(\varphi(\zeta))\right] = \mu(\varphi)/\mu(|\varphi|) > 0,$$

since $\mu(\varphi) \geq \delta > 0$. Furthermore, $\text{sign}(\varphi(\zeta))$ is almost surely valued in $\{-1, +1\}$ since $\zeta \sim \mu_{|\varphi|}$, and so

$$\mathsf{E}\left[\text{sign}(\varphi(\zeta))\right] = p - (1 - p) = 2p - 1,$$

where $p := \mathsf{P}(\text{sign}(\varphi(\zeta)) = 1) > \frac{1}{2}$. From the bound

$$2p - 1 = \mu(\varphi)/\mu(|\varphi|) \geq \delta / \|\varphi\|,$$

it follows that $p \geq \frac{1}{2} + \frac{\delta}{2\|\varphi\|}$. Since $q = \mathsf{P}(\text{sign}(\varphi(\zeta)) = -1) = (1 - p)$, simulating a $(2p - 1)$-coin is equivalent to simulating a $(1 - 2q)$-coin with $q \leq \frac{1}{2} - \frac{\delta}{2\|\varphi\|}$, and this is exactly what $(1 - Y)$ is. It follows that $\mathsf{E}(W) = \mu(\varphi)$ and $\mathsf{P}(0 \leq W \leq \|\varphi\|) = 1$. □

Remark. It is clear that a number of variance reduction techniques can be used to modify Algorithm 11. For example, one could simply average a number of outputs of the algorithm. In addition, one could average different values of $|\varphi(\xi)|$ and/or the $(1 - 2q)$-coins.

An immediate Corollary of Proposition 15 is that more general constrained unbiased estimation problems can be solved.

Corollary 4. *Let φ satisfy $a \leq \varphi(x) \leq c$ and assume we can simulate random variables according to μ. If $\mu(\varphi) \geq \delta > b$ for some known δ then we can simulate a random variable W such that $\mathsf{E}(W) = \mu(\varphi)$ and $\mathsf{P}(b \leq W \leq b + \max\{b - a, c - b\}) = 1$.*

Proof. Let $\tilde{\varphi} = \varphi - b$. Then $\|\tilde{\varphi}\| = \max\{b - a, c - b\}$ and by Proposition 15 we can simulate a random variable $\tilde{W}$ with $\mathsf{E}\tilde{W} = \mu(\tilde{\varphi})$ and $\mathsf{P}(0 \leq \tilde{W} \leq \|\tilde{\varphi}\|) = 1$. We conclude by letting $W = \tilde{W} + b$. □

It follows from Corollary 4 that if we can simulate a p-coin with $p \geq \delta > \epsilon$, then we can simulate a random variable V satisfying $\mathsf{E}(V) = p$ and $\mathsf{P}(\epsilon \leq V \leq 1) = 1$ when $\epsilon \leq 1/2$, by taking μ to be Bernoulli(p), $\varphi(x) = x$, $a = 0$, $b = \epsilon$ and $c = 1$. While we do not recommend the use of Algorithm 11 to obtain a solution to Problem 2, it may be useful more generally. As a simple example, consider the case where μ is the uniform distribution on $[0, 3\pi]$ and $\varphi = \sin$. Then $\mu(\varphi) = 2/(3\pi) > 0$ and one use Algorithm 11 to produce a random variable with expectation $\mu(\varphi)$ and which is almost surely in $[0, 1]$.

References

Andrieu, C., N. Chopin, A. Doucet, and S. Rubenthaler (2012). Perfect simulation for the Feynman–Kac law on the path space. Preprint arXiv:1210.0376.

Andrieu, C., A. Doucet, and R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society B 72*(3), 269–342.

Andrieu, C., A. Lee, and M. Vihola (2013). Uniform ergodicity of the iterated conditional SMC and geometric ergodicity of particle Gibbs samplers. Preprint arXiv:1312.6432.

Andrieu, C. and G. O. Roberts (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics 37*(2), 697–725.

Andrieu, C. and M. Vihola (2012). Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. Preprint arXiv:1210.1484v1.

Asmussen, S., P. W. Glynn, and H. Thorisson (1992). Stationarity detection in the initial transient problem. *ACM Transactions on Modeling and Computer Simulation (TOMACS) 2*(2), 130–157.

Athreya, K. B. and P. Ney (1978). A new approach to the limit theory of recurrent Markov chains. *Transactions of the American Mathematical Society 245*, 493–501.

Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics 164*(3), 1139–1160.

Bednorz, W., K. Łatuszyński, and R. Latała (2008). A regeneration proof of the central limit theorem for uniformly ergodic Markov chains. *Elect. Comm. in Probab 13*, 85–98.

Blanchet, J. and X. Meng (2007). Exact sampling, regeneration and minorization conditions. Preprint, Dept. of Statistics, Harvard University.

Brockwell, A. E. and J. B. Kadane (2005). Identification of regeneration times in MCMC simulation, with application to adaptive schemes. *Journal of Computational and Graphical Statistics 14*(2).

Brown, M., E. A. Peköz, and S. M. Ross (2010). Some results for skip-free random walk. *Probability in the Engineering and Informational Sciences 24*(04), 491–507.

Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Annals of Statistics 32*(6), 2385–2411.

Chopin, N. and S. S. Singh (2013). On the particle Gibbs sampler. Preprint arXiv:1304.1887v1.

David, H. A. and H. N. Nagaraja (2003). *Order statistics* (3rd ed.). NJ: John Wiley & Sons.

Del Moral, P. (2004). *Feynman-Kac formulae: genealogical and interacting particle systems with applications*. Springer Verlag.

Del Moral, P. and A. Doucet (2004). Particle motions in absorbing medium with hard and soft obstacles. *Stochastic Analysis and Applications 22*(5), 1175–1207.

Del Moral, P. and A. Doucet (2014). Particle methods: An introduction with applications. In *ESAIM: Proceedings*, Volume 44, pp. 1–46. EDP Sciences.

Del Moral, P., A. Doucet, and A. Jasra (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society B 68*(3), 411–436.

Del Moral, P. and L. Miclo (2000). Branching and interacting particle systems approximations of Feynman–Kac formulae with applications to non-linear filtering. In J. Azéma, M. Ledoux, M. Émery, and M. Yor (Eds.), *Séminaire de Probabilités XXXIV*, Volume 1729 of *Lecture Notes in Mathematics*, pp. 1–145. Berlin: Springer.

Doucet, A. and A. M. Johansen (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and B. Rozovsky (Eds.), *Oxford Handbook of Nonlinear Filtering*. Oxford University Press.

Eisenberg, B. (2008). On the expectation of the maximum of iid geometric random variables. *Statistics & Probability Letters 78*(2), 135–143.

Flegal, J. M. and R. Herbei (2012). Exact sampling for intractable probability distributions via a Bernoulli factory. *Electronic Journal of Statistics 6*, 10–37.

Foss, S. G. and R. Tweedie (1998). Perfect simulation and backward coupling. *Stochastic models 14*(1-2), 187–203.

Gilks, W. R., G. O. Roberts, and S. K. Sahu (1998). Adaptive Markov chain Monte Carlo through regeneration. *Journal of the American statistical association 93*(443), 1045–1054.

Gordon, N. J., D. J. Salmond, and A. F. M. Smith (1993, apr). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F 140*(2), 107–113.

Henderson, S. G. and P. W. Glynn (2003). Nonexistence of a class of variate generation schemes. *Operations Research Letters 31*(2), 83–89.

Hobert, J. P., G. L. Jones, B. Presnell, and J. S. Rosenthal (2002). On the applicability of regenerative simulation in Markov chain Monte Carlo. *Biometrika 89*(4), 731–743.

Hobert, J. P., G. L. Jones, and C. P. Robert (2006). Using a Markov chain to construct a tractable approximation of an intractable probability distribution. *Scandinavian Journal of Statistics 33*(1), 37–51.

Hobert, J. P. and C. P. Robert (2004). A mixture representation of π with applications in Markov chain Monte Carlo and perfect sampling. *Annals of Applied Probability*, 1295–1305.

Holtz, O., F. Nazarov, and Y. Peres (2011). New coins from old, smoothly. *Constructive Approximation 33*(3), 331–363.

Huber, M. (2014). Nearly optimal Bernoulli factories for linear functions. *Combinatorics, Probability, and Computing*. To appear.

Jacob, P. E. and A. H. Thiery (2013). On non-negative unbiased estimators. Preprint arXiv:1309.6473.

Jarner, S. F. and G. O. Roberts (2002). Polynomial convergence rates of Markov chains. *Annals of Applied Probability*, 224–247.

Keane, M. and G. L. O'Brien (1994). A Bernoulli factory. *ACM Transactions on Modeling and Computer Simulation (TOMACS) 4*(2), 213–219.

Kendall, W. S. (2004). Geometric ergodicity and perfect simulation. *Electronic Communications in Probability 9*, 140–151.

Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics 5*(1), 1–25.

Lamperti, J. (1960). Criteria for the recurrence or transience of stochastic process. i. *Journal of Mathematical Analysis and applications 1*(3), 314–330.

Łatuszyński, K., I. Kosmidis, O. Papaspiliopoulos, and G. O. Roberts (2011). Simulating events of unknown probabilities via reverse time martingales. *Random Structures & Algorithms 38*(4), 441–452.

Lee, A. and K. Łatuszyński (2012). Variance bounding and geometric ergodicity of Markov chain Monte Carlo kernels for approximate Bayesian computation. Preprint arXiv:1210.6703.

Lee, A., C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes (2010). On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *Journal of Computational and Graphical Statistics 19*, 769–789.

Lindsten, F., R. Douc, and E. Moulines (2014). Uniform ergodicity of the particle Gibbs sampler. Preprint arXiv:1401.0683v1.

Meyn, S. P. and R. L. Tweedie (2009). *Markov chains and stochastic stability* (2 ed.). Cambridge University Press.

Mitrophanov, A. Y. (2005). Sensitivity and convergence of uniformly ergodic Markov chains. *Journal of Applied Probability*, 1003–1014.

Møller, J. and R. P. Waagepetersen (2004). *Statistical inference and simulation for spatial point processes.* Chapman and Hall.

Murdoch, D. J. and P. J. Green (1998). Exact sampling from a continuous state space. *Scandinavian Journal of Statistics 25*(3), 483–502.

Mykland, P., L. Tierney, and B. Yu (1995). Regeneration in Markov chain samplers. *Journal of the American Statistical Association 90*(429), 233–241.

Nacu, Ş. and Y. Peres (2005). Fast simulation of new coins from old. *The Annals of Applied Probability 15*(1A), 93–115.

Nummelin, E. (1978). A splitting technique for Harris recurrent Markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 43*(4), 309–318.

Nummelin, E. (1984). *General irreducible Markov chains and non-negative operators*, Volume 83. Cambridge University Press.

Propp, J. G. and D. B. Wilson (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random structures and Algorithms 9*(1-2), 223–252.

Stewart, L. and P. McCarty Jr (1992). Use of Bayesian belief networks to fuse continuous and discrete information for target recognition, tracking, and situation assessment. In *Aerospace Sensing*, pp. 177–185. International Society for Optics and Photonics.

Thomas, A. and J. H. Blanchet (2011). A practical implementation of the Bernoulli factory. Preprint arXiv:1106.2508.

Von Neumann, J. (1951). Various techniques used in connection with random digits. *Applied Math Series 12*(36-38), 1.

Whiteley, N. (2013). Stability properties of some particle filters. *The Annals of Applied Probability 23*(6), 2500–2537.

Whiteley, N. and A. Lee (2014). Twisted particle filters. *The Annals of Statistics 42*(1), 115–141.